Tom Rosamilia bids farewell...

I'd like to inform you that I've accepted a new position as Vice President, Worldwide Data Management Development and General Manager, Silicon Valley Lab in San Jose, California, managing IBM®'s data management products. As I begin my new assignment, I wanted to take a moment to thank everyone that has made my experiences at SHARE so rewarding and productive. In particular, I'd like to thank the SHARE Board, Ribbon Wearers and many customers who have taught me so much over my many years as a SHARE participant.

I wish all of you much continued success. Have one for me at SKIDS!

Tom

Greetings from Paul Loftus...

I'm sure that many of you are thinking 'deja vu' as you read this. Although I am the new Vice President of Systems Software, I am not new to SHARE. In the late 80s and early 90s, I was an active SHARE member while I managed the VM organizations in Kingston, NY and Endicott, NY. Since that time, I've held various positions within corporate IBM. Now I'm back in development, managing both zSeries and iSeries software, and eager to rekindle my past SHARE relationships and establish new ones.

I look forward to seeing many new and familiar faces in Minneapolis in July.

Paul

IBM.

HOT TOPICS

A z/OS Newsletter
Issue 5 August 2001

Contents

In this issue ...

Welcome to our fifth issue! Actually, we're starting to get concerned. Global warming is bad enough without our sizzling HOT TOPICS heating up the atmosphere every six months. Just kidding of course. We've heard lots of good things about how you enjoy reading the articles and finding them useful. And we plan on bringing you many more.

Our issue this time begins, as you've seen, with Tom Rosamilia's bidding us a fond farewell. We wish him the best of luck in his new assignment. And we welcome Paul Loftus, an old friend, to our flagship operating system.

Unlike previous issues, we don't have an overriding 'theme' this time, but what we lose in concentrated focus, we gain in scope, offering articles that range from a discussion of the new architecture to an article that explains how to shop for zSeries on the Internet. Check out the table of contents to get a feel for what's inside.

You'll also note a 'tip-in' — which is a publishing term for a

Kerry Collia

non-magazine type item included with a publication. In our case, it's a business card-sized CD-ROM that contains many Web addresses of sites and pages of interest to you. All you'll need to do is click on the URL.

Because our combined five issues now number at least several hundred pages, we thought it would be a good idea to include an index to help you locate specific articles of interest to you.

In addition, to help us serve your needs better, we've included a short questionnaire. Please take a few moments to fill it out and drop it in your nearest mailbox. To all

those who respond, we will send invitations to an all-expense paid trip to Poughkeepsie later in the year. You will stay at the luxurious downtown Marriott Inn, dine at the exclusive "Chez Pierre" restaurant on steak au poivre, caviar, and Dom Perignon. After dinner, you will enjoy a Hudson river sunset from the deck of Henry Hudson's "Half Moon," where Luciano Pavarotti will sing, in that inimitable way of his, carefully selected renditions of some of your favorite operatic arias. In the morning you will enjoy a sumptuous breakfast in your room, courtesy of room service at the Marriott. Then individual limos will whisk you off to Stewart Airport for a quick trip home in First Class. You will sink into your luxurious leather seat, sighing with contentment and softly saying, hey, those folks at Hot Topics, they're a good bunch!

We are, of course, but we also let our imaginations run wild sometimes. Enjoy.

The Editors

z/OS and 64-Bit: The value proposition, today, and into the future

GREG DYCK

IBM's @server™ z/Architecture is a technically unique and complete definition for a 64-bit execution model, encapsulating and super-setting the existing ESA/390 architecture. It provides 128-bit PSW, 64-bit registers, 64-bit addressing, 64-bit calculations, and virtual address translation support for > 2GB of real and virtual. This architecture was designed to allows z/OS to enter into the 64-bit arena, providing additional value to IBM's customers, while at the same time providing downward compatibility for existing applications. The customer value includes increased storage capacity for constrained systems, improved scalability for faster processors, and migration transparency for applications designed for a 64-bit environment.

Providing exploitation of the full z/Architecture within z/OS is a major work effort. In order to provide value to customers as quickly as possible, the effort has been broken up into a multiple-stage plan that is taking over several years. The IBM @server zSeries 900 processor that z/OS runs on, however, already allow for the future by implementing the full z/Architecture. The z900 is a true 64-bit processor design, with 64-bit calculations taking the same time as 32 bit calculations, and will be able to run future versions of z/OS, such as z/OS 1.2, which exploits additional portions of the 64-bit architecture.

z/OS 1.1 - Here today

z/OS 1.1 (as well as OS/390 2.10) provides the initial 64-bit support by exploiting a subset of the z/Architecture capabilities, including 64-bit registers, 64-bit calculations, and > 2GB of real

Paul Rowntree

storage.[1] This subset implements what is sometimes called "Large Real Support", and allows the z/OS real memory model to scale far beyond the 2GB allowed in ESA/390 mode.

With processors as fast as the zSeries there is an increased opportunity for real storage constraints to be seen with production workloads when real storage is limited to 2GB. Such a constraint manifests itself in the movement of pages to/from expanded and auxiliary storage, and can become a performance bottleneck. Also, real storage constraint increases with increased use of data in memory techniques, as greater amounts of data are kept in real and expanded by data managers to meet performance objectives. Eliminating the 2GB real storage limit generally allows the processor to be utilized more fully and effectively.[2]

A key design point for the large real support was application transparency. Applications which properly use defined general purpose interfaces and services require no change to run in a large real environment. Even interfaces designed to exploit expanded

storage,[3] such as hiperspaces, continue to operate with no change required to the application. When migrating applications to a z/Architecture environment, only applications which did not follow recommended programming practices, such as doing a LRA instruction against unfixed storage or referencing control blocks that were not part of the intended interface, need to be changed. Internal to z/OS, a number of transparent changes have been made to better exploit the large real environment. This includes changes to some access methods to allow their I/O buffer pools to be backed > 2GB when doing I/O. Applications get this value without having to make a single change.

That is not to say, however, that application changes are not sometimes needed to best exploit the large real storage. Making changes to some applications, by both customer and software vendors, to allow fixed storage to be backed > 2GB, specifically those applications that allocate their own I/O buffer pools, will allow them to work better and run faster in a large real environment. And the great thing is, the changes made to interfaces like GETMAIN and STORAGE are downward compatible with OS/390. If you choose to make changes to applications to exploit z/Architecture on a z900, the applications can run unchanged on an OS/390 system running in ESA/390 mode, unless you also used z/Architecture instructions in your change. IBM has made changes to products like IMS™, DB2®, and DFSORT™ to allow them to better utilize real storage > 2GB.

The increase in the maximum amount of real storage provides clear value and constraint relief to customers which have hit the 2GB real storage wall on their system. Some environments will see a significant reduction in system overhead and an improvement in performance by eliminating the need for paging to either auxiliary or expanded storage. It also allows the consolidation of system images that previously had been split apart because they could not be contained in the 2GB of real storage that a system image was limited to. Consolidating these systems can provide value by reducing operations complexity and eliminating duplication of system storage and CPU overhead. The creation of multiple system images has often been required when multiple DB2 or IMS subsystems must be run and the related buffer pools were extremely large in size. Customers now have a choice when confronted with this situation, and instead of splitting the subsystems into multiple system images they may instead move to a z/Architecture environment which does not have the same real storage constraints.

z/OS 1.2 - Coming soon

z/OS 1.2 will provide the initial basic 64-bit virtual storage management support, enabling the growth of private storage to amounts that are almost inconceivable: 8 billion times the current 2GB limit. This support allows applications which are memory intensive to extend their data model vertically beyond 2GB without having to use mechanisms like multiple dataspaces. The storage > 2GB can be used only for data at this time, similar to dataspaces; programs cannot execute in virtual > 2GB.

Management of virtual storage > 2GB will be done using new interfaces and in large increments,

called memory objects. Memory objects are not managed in small 8 byte increments like GETMAIN or STORAGE. The minimum size of a memory object will be 1MB, and it must be allocated in MB multiples.

Basic services required to effectively use the memory objects for buffer pools, and the like, will be supplied in z/OS 1.2, along with diagnostic support. This includes the ability to do I/O to the storage using EXCP, to transfer data to and from the memory object to the coupling facility, and to capture the memory object in SVC dumps.

Application transparency continues to be important. Only applications which wish to exploit virtual storage > 2GB must be changed. The initial value for this support will be in IBM and vendor written middleware exploitation, since significant application changes are required to use it. Also, few customer applications require this much virtual storage. Ideal candidates to exploit virtual storage > 2GB are database managers, such as DB2.

The near future

Following on the foundation laid by z/OS 1.2, additional support will be provided for assembler applications that wish to run in 64-bit addressing mode and exploit memory objects.

Support will be added to allow programs to receive initial control in 64-bit addressing mode (AMODE 64). However, these programs will continue to be loaded in virtual storage < 2GB. Execution of instructions contained in memory objects will still not be supported.

To make the writing of AMODE 64 applications easier, a specific set of system services will be enabled so that they can be

invoked without requiring the application to switch into 31 bit addressing mode first. This enablement does not necessarily change the requirements of the services for where their parameter lists are located. Parameter lists may still be required to be located < 2GB.

The AMODE 64 support and the enablement of z/OS services to be issued in AMODE 64 provide a jumping off point for increased assembler language and middleware exploitation of the 64-bit virtual environment.

The future

In the future, shared memory objects will be provided. With the creation of shared memory objects the basic infrastructure required for high level language environments is in place, and can also be provided.

The shared memory object support will allow assembler applications in different address spaces to share data and communicate using virtual storage > 2GB. This support is the foundation for higher level memory sharing concepts such as those provided by z/OS UNIX System Services.

The high level language environment provided by UNIX® System Services and the Language Environment® run time and class libraries will be enhanced to enable 64-bit application exploitation of memory objects. The high level language support will be targeted specifically for C/C++ applications running in a UNIX System Services environment.

The C/C++ compiler will be extended to generate 64-bit enabled programs using z/Architecture instructions, and to exploit the new 64-bit execution environment provided by UNIX System Services and the Language

Environment. Additionally, 64-bit support will also be provided in the TCP/IP and diagnostic/debug tools.

At this point the 64-bit infrastructure is enabled for the general application environment, and can be utilized in environments like Websphere. It also addresses some of the application porting issues, with applications designed for a 64-bit environment like JAVA® and WebSphere® Application Server.

Way out there

The final piece of the z/Architecture to implement, and the piece least defined today, will probably be allowing program execution in storage > 2GB. As more applications environments move to exploit virtual > 2GB for their data, this will become required. Especially for environments such as Java which dynamically build executable code streams.

While the z/Architecture defines, and hardware always uses, the full 128-bit PSW, z/OS has only tolerated them until now. During interrupt processing, prior to this support, the system software always compressed the 128-bit interrupt PSW into 64-bits for use throughout the system to allow for compatibility. If an interrupt occurred with an instruction address > 2GB it was considered an error, and a failure resulted. With > 2GB execution support this is no longer true.

Execution > 2GB is the last piece of z/Architecture to implement in z/OS, due to the number of control blocks (the RB control blocks in particular) referenced by non-privileged and privileged applications alike that contain 64-bit PSW values. Replacing the 64-bit PSW values with 128-bit PSWs will require introducing some application incompatibilities into this level of z/OS. The trick will be in maintaining some level of backward compatibility for the greater portion of the applications that do not exploit execution > 2GB, yet growing the PSW size for those that do. For non-privileged applications that only look at the PSW in the RB this will probably be done by provided a 64-bit shadow copy of the PSW in the RB. For privileged applications that update the RB control block, however, changes will be required because changes to the shadow copy of the PSW will be ignored. To address this issue, the IEARBUP service was provided in

OS/390 2.10. Applications which exploit IEARBUP to update the RB PSW will not require changes. Applications that **do** exploit execution > 2GB and look at the PSW will have to properly exploit the new interfaces to access the 128-bit PSW.

Conclusion

z/OS is an operating system with a long tradition of compatibility, full of exciting potential, and has a bright future with lots of innovative changes to come. These attributes are all part of the z/OS value proposition, showing it's ability to evolve and meet future needs without abandoning the past. The large real support in z/OS 1.1 is just the first step, but an important step, in the enablement of z/Architecture for z/OS, with its capabilities providing immediate value. The evolution and exploitation of the large virtual capabilities in subsequent releases will continue to provide new value to customers, software vendors, and IBM into the future.

Additional insight into future z/OS 64-bit support can be found in the white paper titled "IBM @server zSeries 900 64-bit Virtual Storage Roadmap" at **http://www-1.ibm.com/servers/ eserver/zseries/library/ whitepapers/gm130076.html** •

1. A software implementation limit of 128GB for real memory exists at this time.
2. Some environments that are not real storage constrained might see a small reduction in performance; for more information see *ftp://ftp.software.ibm.com/ software/mktsupport/techdocs/64bit.pdf*
3. Support for expanded storage was removed from z/OS in z/Architecture mode, since real storage is no longer limited to 2GB and real storage > 2GB can take the functional place of expanded storage.

Expanded use of z/OS Managed System Infrastructure (msys) for Setup

FRANZ-PETER BOLEY

In the last issue of z/OS Hot Topics, a new technology for configuring z/OS was described. This technology – z/OS Managed System Infrastructure (msys) for Setup – is delivered as part of z/OS Release 1. And more support is planned in z/OS Release 2.
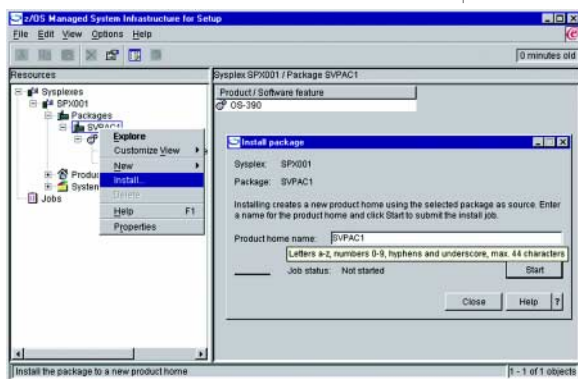
"Smart" configuration dialogs

With z/OS Release 1, msys for Setup introduces easy-to-use workstation-based dialogs for specifying a configuration. The dialogs can understand the current system configuration and take it into account. An automated process then updates the system with the defined configuration. All this is centered around a z/OS Management Directory, which becomes the central location for all configuration values.

In z/OS R1, the msys for Setup framework is used by Parallel Sysplex® functions to set up a Parallel Sysplex resource sharing environment.

Customers like it

Customers who have seen z/OS msys for Setup so far are excited about this new technology. z/OS has made a big step to make the S/390 environment more easily

manageable. Many customers are concerned about the lack of system programmer skill in the next few years, because many S/390 system programmers will retire in the not too distant future. It is quite difficult (and expensive) to fill the slots appropriately. msys for Setup addresses these issue by reducing the skill requirements needed for configuring z/OS and z/OS products and making experienced system programmers more productive. Configuration activities that typically take a week when done manually can now be done within a day with msys for Setup.

Extending the value of msys for Setup

With z/OS Release 2, additional z/OS components will use msys for Setup for their configuration, extending msys for Setup's direct value to your IT staff through its productivity enhancements and reduced skill requirements.

- **TCP/IP Services** will use msys for Setup for defining basic TCP/IP settings, for an easy definition of network devices and links, and for the setup of an FTP or Telnet 3270 server. After customization, the corresponding TCP/IP data sets (like tcp.data, profile.tcp, and omproute) are created automatically. Customers who are new to TCP/IP will especially benefit from these easy-to-use configuration dialogs.
- **Parallel Sysplex** will use msys for Setup to cover additional tasks, such as the setup of the system logger, which includes

the logger requirements for IBM License Manager, for OPERLOG, and for LOGREC, as well as defining the options for enhanced catalog sharing. As in z/OS Release 1, msys for Setup will understand any existing definitions, so that they can be easily managed using the msys for Setup dialogs.
- With z/OS R2 you will also be able to set up a **Base Sysplex** environment using msys for Setup. It will have all functions of the Parallel Sysplex support relevant for Base Sysplex, including the support for migrating from a Base Sysplex to a Parallel Sysplex.
- **ISPF** will use msys for Setup to create the ISPF configuration table keyword file and load module.

With msys for Setup, z/OS is building ease-of-use right into its platform. The msys for Setup framework comes as an integral part of z/OS, and the new support mentioned here is shipped as part of the corresponding z/OS components. With no extra charge, you can benefit from the value of this new and exciting technology.

Additional information

If you missed the msys for Setup article in z/OS Hot Topics Issue 4, it's available at: **http://publibz.boulder.ibm.com/epubs/pdf/e0z2n100.pdf**

For more information on z/OS msys for Setup, go to:
http://www.ibm.com/eserver/zseries/msys

Or contact:
Uwe Gramm
z/OS Planning
gramm@de.ibm.com •

Linux and USS: Heavy hitters for zSeries

ROBERT NAUMANN AND J.D. ROSS

In the IBM @server world, there is a new player: a new, exciting, lean, mean technology that has the ability to change the playing field that today's server revolves around. Linux® has arrived, and gained the attention and focus of IBM, along with the entire business world.

Linux is the fastest growing implementation of a UNIX based operating system in the server market. In fact, with International Data Corp. reporting Linux having 132% growth in the year 2000, it is the *fastest growing server operating system produced!* So what makes IBM's adoption of Linux so special? Well, for the first time, developers have the opportunity to bring their applications and tools rooted in UNIX technology to the legendary reliability of enterprise hardware manufactured by IBM. The door is now open to new and powerful uses for the mainframe.

Linux—paired with over three decades of innovation and ingenuity that comprise z/OS—makes the zSeries a potential host to all workloads, new and existing. The two operating systems work seamlessly alongside one another, allowing a business to leverage the strengths of each operating system on the same machine. IBM has embraced the differences between z/OS and Linux, giving developers the freedom to choose which operating system will best meet the needs of their application. For a developer, an examination of the Qualities of Service offered by Linux and z/OS is needed to make good decisions as to how each workload will be handled efficiently and effectively.

Linux on the zSeries provides new developers a familiar path to

the mainframe landscape. Boasting fast development time, low development costs, a flexible application environment, and the backing of IBM, Linux is making strong headway into server environ-

Tracy Bondi

ments all around the world. IBM is expediting this adoption with a commitment to invest over one-billion dollars in Linux development, service, and support in 2001. A major advantage in the effort to bring Linux to the mainframe is the ease with which existing Linux code, originally written for Intel x86-based architecture or other hardware platforms, can be ported to the IBM @server family. Further enhancing this portability, the top distributors such as SuSE,® TurboLinux,® and Red Hat® have made their distributions of Linux compatible with the zSeries platform.

UNIX System Services is essential to the e-business run time environment of z/OS. It is a key element in IBM's support for middleware and applications, and is an integral part of some of IBM's most strategic solutions such as WebSphere EE. It provides a sound and robust UNIX environment, carrying with it the qualities of service (i.e., security, availability, scalability) expected and

delivered with z/OS. IBM will continue to enhance USS to support the e-business™ infrastructure, as well as for vendor and customer applications running on z/OS.

While there are dozens of qualities on which a comparison between the two operating systems could be based, five main categories remain paramount in the eyes of the enterprise customer. They are: Performance, Portability, Security, Availability, and Scalability.

Performance

Linux provides a thin, flexible OS with the ability to be tailored for an application-specific use. Developers have the freedom to implement services found under USS (i.e. address space isolation, memory management, and classification of work for workload and resource management) themselves. This "custom-fit", with its decreased complexity, can result in increased application performance when compared to USS, as well as giving the developer the ability to implement only the services necessary to meet their requirements.

USS can provide an application with the tools necessary to perform well by letting the operating system handle tasks such as memory management at the operating system level. While not allowing application-tailored API's, USS provides refined, and tested services so developers can realize their performance requirements without having to devote development resources to the back-end functionality.

Portability

Linux gives developers a quick and efficient way to bring applications to the zSeries platform. Applications will instantly benefit from the Hardware Qualities of Service offered by zSeries and have a starting block from which working ideas can demonstrate proof-of-concept, with a goal of finally being transformed into full-fledged z/OS programs. In a server environment where short-term deadlines must be met, Linux makes an easy transition to the zSeries mainframe.

USS allows applications to benefit from the full spectrum of Qualities of Service offered by z/OS and the zSeries. The initial re-architecture of the code will be met with less hesitation once an idea is proven to be workable on a Linux image. Once running on z/OS, the application is a robust, business ready tool with the ability to natively exploit both the zSeries and any legacy data housed therein. z/OS and USS gives developers a long-term solution for porting from other platforms to the proven reliability of the zSeries.

Security

Linux provides developers with a multithreading, multi-user operating system capable of isolating critical tasks or processes through the use of VIF (Virtual Image Facility) and VM (Virtual Machine). Pairing a single application to a single Linux image allows a developer to assure many of the levels of data integrity expected by users of the zSeries. Through the use of technologies such as LDAP (Lightweight Directory Access Protocol), multiple Linux images have access to a shared repository of user privileges and access permissions.

USS allows programmers new to the zSeries platform to work using the UNIX-Security model. This familiar environment can help the transition from a UNIX world to the enterprise mainframe. Developers can take advantage of all z/OS services and API's, such as address space isolation, to provide their application with the level of security inherent to the operating system. Data integrity is paramount in z/OS, and developers will be able to exploit this over multiple applications without the need of managing multiple images.

Availability

Linux boasts heightened availability status through attributes found in hardware, on the zSeries, alone. The zSeries has matured into a "self-healing" machine, boasting features such as the ability to replace faulty processors and memory with working spares on-the-fly, with no manual intervention. However, a developer must contend not only with hardware failure, but software failure as well. Right now, Linux images must be implemented with high-availability clustering software such as the package shipped with TurboLinux for Intel that enable fast drop-and-recover of Linux images.

USS on z/OS has the built-in resiliency to ensure that the operating system, and the applications running, can recover in the event of a fault. Having processes built into the operating system decreases the dependency on clustering currently experienced by Linux. Coupled with IBM's Sysplex clustering technology, z/OS provides the developer a rock-solid and dependable platform to deploy multiple simultaneous applications.

Scalability

Linux enables quick and relatively easy horizontal growth by the ability of adding more images as workload increases. Applications such as a web-server will benefit from this model without much tuning of the individual image. In this way, increased performance in times of demand can be realized using a fast and modular method of expansion. However, Linux (in its current state) is unable to scale vertically in a large SMP environment efficiently, leaving no other option then to rely on multiple images.

Developing for USS allows an application to take advantage of the vertical scalability of z/OS. In a hardware environment consisting of 8 or more processors, z/OS can easily and efficiently scale to meet demand. Also, Sysplex clustering allows z/OS to scale horizontally when the workload requirements are such.

Two powerful weapons

Linux and z/OS are very different, both in the technology they employ and in the business solutions they provide. Linux is not a replacement for z/OS or USS. It is a way to help bring new and different workloads to the IBM mainframe. It gives an IBM developer the power to use the IBM eServer in more ways than ever, consolidating dozens of UNIX or Linux machines onto the best UNIX box ever produced. Linux on the zSeries is the perfect tool for an IBM software developer to demonstrate Proof of Concept projects. Development tools such as Cprof, LTrace, and YAMD allow fast deployment in time-critical situations. These tools, and others, are available on the Linux Application Tools website at: **http://www.ibm.com/servers/ eserver/zseries/os/linux/ldt/**

Once the application has proven itself—if it requires the full qualities of service— it can be ported to the most secure, scalable,

and available operating system to ever grace the big iron - z/OS. Working together, Linux and z/OS broaden the strength and scope of software for IBM's @server. And IBM is committed to working with customers who are looking to bring new workload to the enterprise

server, either through Linux or through USS.

So, when the next project arrives at your desk, you don't have to pick your weapon—Linux or z/OS—you can use both. Workloads, big and small, don't stand a chance. •

The following people contributed their technical expertise to this article:
Mythili Bobak
Frank J. DeGilio
John Eilert
Chris Rohrbach

The hitchhiker's guide to the world of WebSphere for z/OS

PAUL DIMARZIO

This article is an introduction to WebSphere for z/OS for the uninitiated, especially those with a "z" background. If you've heard that WebSphere for z/OS is "hot" but want to know why; if you know something about Java but only in the context of flaming logos and tailored Web pages; or if you're just downright curious, read on.

You are a professional in the zSeries (S/390) community. What if I were to tell you that WebSphere for z/OS opens up the z/OS platform to applications developed by anyone with enterprise Java skills, regardless of their level of z/OS expertise? You would no doubt be skeptical, but hopefully after reading this article, you will see that this is precisely what WebSphere for z/OS brings to the table. To net it out, WebSphere for z/OS is developer technology that, in effect, "takes the mainframe out of mainframe programming."

Market perspective: Application servers

WebSphere for z/OS as a brand encompasses a wide variety of Web infrastructure software products that help companies at each stage of e-business development: from startup, to integrating and exploit-

ing business processes, to handling high-volume Web transactions. The heart of WebSphere for z/OS—called the Foundation—is infrastructure that is commonly referred to as an "application server" in the industry. The application server facilitates the rapid development and deployment of e-business applications by handling transactions and extending more "traditional" business data and applications (like DB2, CICS® and IMS) to the Web.

Note: The use of the term "application server" is different from the popular IT meaning; namely, a server box that is dedicated to running business applications.

Let's dig a little deeper. Since all e-business applications share a common need for a base set of system services (such as persistence, security and transactions), it is readily apparent that the productivity of e-business application developers would be greatly improved if these services could be standardized and presented to applications in a transparent, platform-independent fashion. You will often see such services referred to as "plumbing" by business programmers because they are programming activities that are necessary, yet have no

direct correlation to business processes. The connotation is deliberately negative.

The ultimate purpose of an application server is to hide this plumbing by transparently extending a set of essential system services to e-business applications so that business programmers can easily assemble these applications from compact pieces of business logic without requiring a deep understanding of platform-specific infrastructure. In this rapidly changing world, business programmers cannot afford the time to figure out how to use system-level services to achieve business objectives. This is especially true if business applications are to be deployed across a wide variety of platforms, each with its own model for accessing system services.

The primary value proposition of application servers, therefore, is to facilitate the development of e-business applications by letting business programmers focus solely on writing business logic; the application server handles all of the difficult system-related chores transparently. This is achieved through a programming model that partitions the task of *developing* applications (which does not

require knowledge of platform infrastructure) from the task of *deploying* applications (which does require platform skills). Application servers are not necessarily aimed at easing the deployment of these applications (i.e. If you are in I/T, this technology will not necessarily make your job any easier); the major appeal of this class of middleware is to the people responsible for writing business applications.

Businesses looking to become e-businesses desperately need to reduce the amount of system-level skills needed in order to build their e-business applications. The application server class of middleware plays a key role in the hosting of new Web-oriented applications because it simplifies application development.

The WebSphere for z/OS Foundation application servers support a programming model based on enterprise Java technologies, so let's see how Java fits into this picture and explore its appeal to e-business application developers.

Technology perspective: Java 2, Enterprise Edition

Java®'s early popularity centered around the "Write once, run anywhere" philosophy of application development. An application written in Java could be run on any platform that supported the Java Virtual Machine (JVM). JVMs quickly appeared on all major platforms, and true application portability became a reality for perhaps the first time ever. However, in these early days, Java technology was used almost exclusively for visual, client-side applications because performance was lackluster and Java technology did not support the basic system-level services required by enterprise-class applications. Perhaps most important of all, Java technol-

ogy did little to address the complexity of writing multi-tier, heterogeneous e-business applications.

Java technology has grown up in a hurry, and enterprise Java now provides the most pragmatic approach to e-business application development in the industry. The key to the enterprise Java programming model is the concept of container-based components.

A container is a software runtime environment that provides common services on behalf of the Java components that run in the container. The container-based component model allows the developer to defer the execution of critical computing tasks—such as maintaining a secure environment, executing transactions and persisting data to a resource manager—to the runtime containers. A business programmer merely specifies what services are required by an enterprise Java component (using XML-based tags called deployment descriptors), leaving the implementation of these services to the container. The container also exposes a platform-independent connector architecture that allows non-Java data and applications to be seamlessly integrated with new Java business logic. No plumbing required! Enterprise Java technology provides an environment in which business programmers with Web-oriented skills can produce meaningful, advanced enterprise applications without having platform-specific coding expertise.

The Enterprise Java programming model defines three container-driven Java components:
- **Java Server Pages (JSPs),** to handle server-side presentation logic.
- **Java Servlets,** which are ideal for controlling the flow of execution.

An animation's worth a thousand words
MARY ELLEN COLEMAN

Need a gentle introduction to WebSphere for z/OS and how you can use it in your business? Check out our new animation at: *http://www.ibm.com/software/webservers/appserv/cb/wasanims.html*

We've created three new short pieces that describe the application server, its characteristics, and its business value. All are available in Macromedia® Shockwave™ format. We also include the scripts in Adobe® Acrobat® format.

As you develop applications, you focus on business requirements, and delegate supporting services like security, component location,
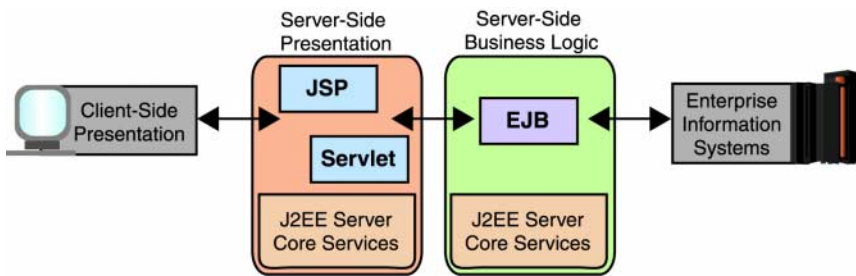
The animation uses graphics and narration to explain the concepts that are difficult to illustrate in printed documentation. We think they're fun! Please let us know if they help you and if there are additional topics you would like to see animated.

- **Enterprise JavaBeans (EJBs),** the powerful object-oriented technology for representing business logic and facilitating the access to non-Java systems and data.

Combined with a set of useful services and connectors for direct access to back-end systems, these container-driven Java components make up the backbone of Java 2 Enterprise Edition (J2EE), the industry's premiere model for building e-business applications.

The use of J2EE containers very clearly separates the roles of business application developer, container provider, and application installer. It allows the developer to use a platform-neutral, workstation-based application development tool to build applications that can be deployed on any J2EE server—including z/OS—by deferring the resolution of platform-specific details to deployment time.

Reality perspective: Not all application servers are equal

Although all Java-based application servers should, in theory, provide a uniform programming model to the developer, they will vary greatly in their ability to provide Enterprise-class qualities of service on behalf of the applications that they host. Hardware (IBM @server zSeries 900, S/390®, Intel®, RISC, etc.) and the operating system (z/OS, OS/390, Microsoft Windows®, UNIX, LINUX, etc.) have always been

important decision factors in choosing where to host an application. With the J2EE container-based model for developing applications, choosing the right application server becomes the predominant development decision point, and choosing a hardware/software platform is deferred to deployment time (and based on the quality of service required).

Consider that, in more traditional methods of developing applications, the programming

model required that access to system services be explicitly coded into the application. The developer was responsible for, invoking the transaction service, establishing security protocols, utilizing efficient threading techniques, and much more. Once a hardware/OS platform was chosen, a skilled programmer could create highly-efficient code that took full advantage of the quality of service offered by the chosen platform.

With J2EE, the responsibility for efficient use of system resources shifts from the developer of the application to the provider of the container. Business logic is the same on all platforms; what differ from one platform to another are things like data access, transactions and how one provides security. All of these capabilities are provided by the container, and the degree to which an application server vendor can provide J2EE containers that exhibit the required qualities of service can make or break an

e-business application. Some new factors that become important decision points include:

- **Availability on multiple platforms.** Even the most efficient application server is limited to providing qualities of service that can be supported by the underlying platform. To meet the most diverse set of requirements, an application server should be available on a wide variety of platforms (and, of course, be fully optimized to each).
- **Ability to leverage existing applications and data.** In the large enterprise space, the ability to incorporate existing resources into new J2EE applications is critical.
- **Available tools for installation.** Although application servers present a neutral platform view to the developer, personnel responsible for application deployment will likely need to learn new skills to install applications in a particular server environment. A key point of differentiation will be the quality of the tooling that is provided to ease the task of deploying applications.

In order to successfully offer a suite of Enterprise Java application servers that meet these requirements, a vendor must have in-depth knowledge and experience with all of today's top hardware/OS platforms; an equally strong background in the major transaction monitors, data managers and messaging systems, and an integrated suite of development and management tools to leverage into this environment. IBM excels in all these areas, and has put that expertise to work in its WebSphere for z/OS family of application servers.

## IBM's J2EE foundation: WebSphere application servers

Hopefully you now realize the business value of application server middleware to the development of new business applications, and how an application server built on the J2EE programming model "takes the mainframe out of mainframe programming" by shifting platform-specific responsibilities from the application to the container. Given this shift of responsibility, it should be evident that the implementation of the application server itself is critical to the overall success of business applications.

The WebSphere Application Server V4.0 for z/OS and OS/390 extends the zSeries role as a premier server for enterprise e-business, and fully exploits the z/OS hallmark qualities of service such as scalability, performance, security, and availability. Even though applications are developed on a workstation by programmers unfamiliar with z/OS, these applications—when deployed in WebSphere for z/OS containers on z/OS—inherit the qualities of service expected of the "z" platform. That means, for example, that J2EE applications deployed on z/OS automatically enjoy the benefits of Workload Manager and Parallel Sysplex for scale and availability. Rapid creation of production-ready e-business applications, integrated with the extensive inventory of existing z/OS applications and data—that's the WebSphere for z/OS story.

To sum it up, the WebSphere for z/OS Foundation implements an industry-standard programming model around J2EE that lets you build platform-independent applications and deploy them virtually anywhere. When these applications are deployed on z/OS, they automatically inherit the qualities of service that z/OS has to offer.

Beyond the foundation
Rounding out the WebSphere for z/OS brand—and resting on the Foundation application servers—is a set of e-business infrastructure products designed to meet a wide range of needs:

- **Development extensions:** Provide the full range of tools needed for e-business, bridging Java programmers, graphic designers, and Web designers.
- **Presentation extensions:** Provide the ability to manage the user experience at an individual, personalized level, as well as extend that experience beyond the Web, utilizing voice media and pervasive devices.

- **Deployment extensions:** Provide e-business scalability to fit any size or complexity of e-business solution, including performance management, site management, and site analysis.
- **Integration extensions:** Provide business process automation and workflow management, as well as the ability to integrate with backend systems.
- **Accelerators:** Modular and extensible business services that enable companies to anticipate and react swiftly to business demands and opportunities, while differentiating their business practices and deliverables. •
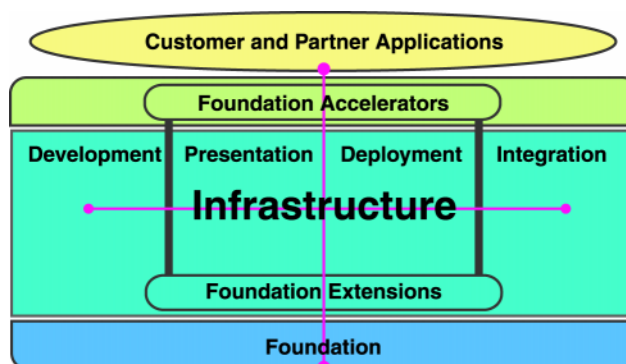
WebSphere for z/OS customization that fits like a glove!

CHRIS HASTINGS AND KIRSTEN MCDONALD

We've reduced the time and effort required to install and customize WebSphere Application Server V4.0 for z/OS and OS/390 (WebSphere for z/OS). Similar to the z/OS Web-based wizards, the dialog guides you through defining WebSphere for z/OS customization options. The dialog generates jobs and samples according to options you supply, then builds a tailored set of instructions that take you through your initial WebSphere for z/OS setup.
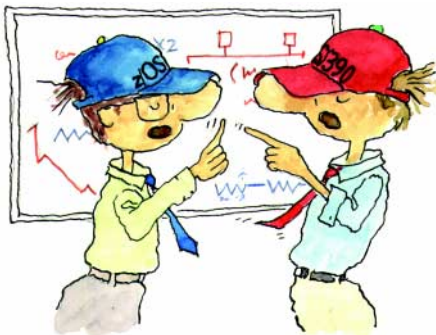
Paul Rowntree

The dialog eliminates the need to hand-tailor sample jobs supplied with the product. You define the customization options once in the dialog panels, then the dialog generates the jobs with your options, eliminating the need to define them in several places. The benefit to you is reduced typos and inconsistencies, and a quicker customization.

Experts at IBM's Washington Systems Center prototyped the customization dialog. Their expertise helped give the dialog real-world applicability. After prototyping, a team of code developers, usability experts, testers, and information developers completed the development and testing work.

The Customization Dialog will be delivered as a WebSphere for z/OS V4.0 service update. Watch the service stream for APAR PQ48858.

The dialog is one part of the overall installation and customization process

The customization dialog is intended for the system programmer or administrator responsible for installing and customizing WebSphere for z/OS. The dialog is intended to be used only once, for the first time you customize the product. Please see the chart below for how the customization dialog fits into the overall installation and customization process.

When you finish the entire customization process, you have WebSphere for z/OS running in a monoplex system. As you gain experience, you can roll out WebSphere for z/OS across your sysplex to gain the advantages of z/OS sysplex operations.

As you can see, the dialog covers a portion of WebSphere for z/OS customization. Specifically, it creates tailored jobs to:
- Define the WebSphere for z/OS run-time configuration (systems management server, naming server, interface repository server, daemon server).
- Create the system management HFS structure and the initial environment file.
- Create and customize the LDAP server.

Stage	Description
1	Install WebSphere for z/OS V4.0 using SMP/E according to the Program Directory.
2	Install prerequisite products, such as DB2 for z/OS Version 7. Configure z/OS subsystems, such as resource recovery services (RRS) and workload management.
3	Run the customization dialog. Through a series of panels, you choose options and define variables. Using your values, the dialog tailors the WebSphere for z/OS customization jobs but does not execute them. Rather, the dialog provides a custom set of instructions for you to follow. When you finish the dialog, you have a set of instructions and tailored jobs ready to complete the product customization.
4	Follow the instructions created by the customization dialog. The instructions take you through a process we call the bootstrap. When you finish, you have a complete WebSphere for z/OS run-time configuration.
5	Install the workstation-based Administration and Operations applications (also called the System Management User Interface). With the Administration application, create the J2EE or CORBA (MOFW) server definitions. These servers are used by the installation verification program (IVP) and are examples of application servers you will create for your own applications.
6	Run the J2EE or CORBA (MOFW) IVPs to verify your WebSphere for z/OS system is working properly. You're done!

- Run the installation verification programs (IVPs).
- Set up WebSphere for z/OS security controls (RACF®).
- Copy the generated jobs into your system libraries.

Using the customization dialog

To start the dialog, you execute the following command in ISPF:

ex 'bbo.sbboclib(bbowstrt)'
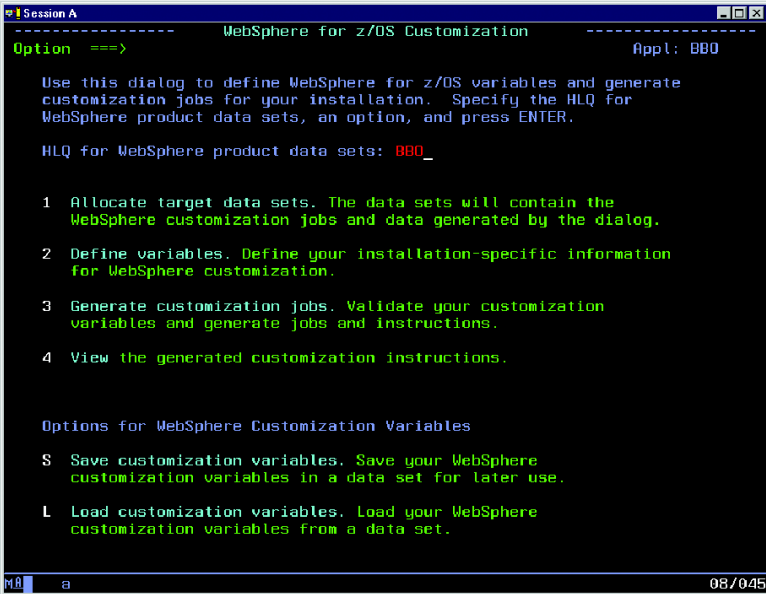
After the system displays a splash screen and you press **Enter,** you see the main menu (see graphic above).

First, you allocate the data sets to contain the customized jobs and instructions (Option 1).

Then, you define the variables for your system (Option 2). If you have not already loaded initial variables, the dialog asks you to prime the customization variables first. If you have already loaded the variables, you see the Define variables menu. The defaults for the variables are shown on the panels. In most cases, the defaults are acceptable. There are a few you must supply, such as WebSphere for z/OS IP information, the system

name, and the sysplex name. Help information is available for each panel by pressing PF1.

Once the variables are defined, you generate the jobs (Option 3). This step validates the variables defined and generates the tailored jobs. If the dialog finds any validation errors, it generates a report. You can either return to the

Define variables option to correct the errors, or you can continue to generate the jobs.

Once you have generated the jobs, you can view (and, if you want, print) the generated instructions (Option 4).

You can save your customization variables in a data set for later use if you do not complete your definitions in one sitting (Option S). To restore variables you previously defined, use the load the variables option (Option L).

For more information about WebSphere for z/OS customization, see *WebSphere Application Server V4.0 for z/OS and OS/390: Installation and Customization,* available online at:
http://www.ibm.com/software/ webservers/appserv/

Using z/OS wizards for base system requirements

You may find the z/OS wizards useful in customizing the WebSphere for z/OS prerequisites. The diagram below shows the wizards that might help. •

WebSphere for z/OS requirement for:	Wizards to help:
• UNIX System Services	• UNIX System Services Configuration Wizard
• TCP/IP	• TCP/IP Configuration Wizard
• Log streams	• Parallel Sysplex Configuration Wizard
• RRS configuration	• Base Sysplex Configuration Wizard

z/OS wizards: A closer look at the z/OS UNIX Customization Wizard

MARILYN ZEPPETELLI

What is a z/OS wizard, anyway?

Most of you are familiar with wizards from the personal computer arena, where many applications use wizards to guide users through such tasks as installing a product, creating a new document, or creating a presentation. While wizards are not usually associated with the mainframe platform, we do provide the z/OS wizards, Web-based tools designed to simplify some of the complex tasks associated with installing and customizing z/OS. Unlike wizards on the PC platform, however, z/OS wizards do not directly change anything on the z/OS system because they are not directly connected to it. To use wizard output on your z/OS system, you need to download the output from the Web, then upload it to your z/OS system.

We offer two types of wizards—planning wizards and customization wizards. **Planning wizards** provide a customized view of documentation based on answers you provide to interview questions. Planning wizards give you access to the information you need most and hide the information that you do not need. **Customization wizards,** on the other hand, provide tailored jobs, commands and information to help you complete the customization task. The output uses 'best practice' values and default values wherever appropriate. Customization wizards gather more information from you, so the customization wizards allow you to save the data you entered in case you don't complete all the interview topics in one sitting. You can return at another time, reload your data, and pick up where you left

Kerry Collia

off. You can use customization wizards both to create a new configuration when you are installing a product for the first time, and to verify an existing configuration by comparing the values the wizard generates to the values you came up with yourself.

Beginning with z/OS V1R1, there is a new element called **msys for Setup,** which provides a graphical user interface for customizing z/OS through wizard-like interviews. In contrast to the Web-based wizards, **msys for Setup,** exploiters are able to actually execute the changes and updates that you make because the exploiters are connected to your z/OS system. The Parallel Sysplex wizard, available with z/OS V1R1, is the first **msys for Setup,** exploiter. Over time, most z/OS customization wizards will probably become **msys for Setup,** exploiters. In the meantime, Web-based wizards let us provide you assistance with these tasks sooner because the wizards are not tied to a release cycle. We can put the wizards out on the Web for you to use as soon as we have coded and tested them.

How do the wizards work?

On the introduction page of each of our wizards, you will find our 'wizard guy' juggling a

different object for each wizard. The "read me first" file provides some background information on what that particular wizard is all about and general comments about how it works. You can start using any wizard by clicking **Start.** You'll go through several different interview topics that ask you for the information that the wizard needs to create your customized output. Once you've answered all the questions with valid data, press **Build** to see your output. The output, presented in a checklist format, takes you through the steps you need to perform to complete the task at hand. A very short survey at the end of every wizard encourages you to tell us what you think of the wizard.

A closer look at the z/OS UNIX Customization Wizard

The z/OS UNIX Customization Wizard is actually two separate wizards in one—a single-system wizard to help you customize z/OS UNIX for the first time, and a shared HFS wizard for implementing a shared HFS environment in a parallel sysplex.

UNIX single-system wizard

The UNIX single-system wizard has four interview topics—"Your System Environment," " z/OS UNIX Processing Values," "File Systems," and "Initial Security Setup." After answering all the interview questions, click on the **Build** button to generate your customized output. (If you answer only the required questions, the wizard will let you build output, but it is likely to be incomplete.) The UNIX single-system wizard generates some parmlib members, JCL jobs, and commands you

Spreading the love: SMP/E for z/OS and OS/390

KURT QUACKENBUSH

In October of this year, many enhancements will be available in SMP/E to simplify and extend the installation and service of software for z/OS and OS/390® systems. There are improvements in the area of SMPPTS management, HOLDDATA research, Hierarchical File System (HFS) management, and you'll even be able to receive input from a network location in addition to traditional tape or DASD inputs.

SMP/E for z/OS and OS/390

The first piece of exciting news is that the next release of SMP/E will be available for free to all OS/390 and z/OS Release 1 licensees. Of course it will also be available as a base element of z/OS Release 2. Previously, SMP/E enhancements could be used only by those on the latest releases of z/OS or OS/390. Now, both you and your products and elements on the platform can take advantage of new installation and packaging functions, regardless of your z/OS or OS/390 level. By the way, it also allows IBM to more quickly exploit some new technologies without waiting for users to migrate to the latest release of z/OS (more on that in the article *"SMP/E: It's not just for tapes anymore!"*.)

Managing the SMPPTS data set

We all know the SMPPTS data set can become extremely large, so much so that a single partitioned data set (PDS) on a single volume is sometimes not big enough. The SMPPTS is used during SMP/E RECEIVE processing as storage for SYSMODs of all types, but, for the most part, it gets filled with PTFs.

You've likely been dealing with your SMPPTS in a number of

ways, such as using the ACCEPTCHECK processing of the RECEIVE command, and REJECT PURGE, but have you tried using more than one SMPPTS data set? *"What?! More than one data set?"* Absolutely... read on.

With the latest release of SMP/E, you can define **SMPPTS spill data sets.** SMPPTS spill data sets contain the overflow from the primary SMPPTS data set. In other words, if the SMPPTS becomes full, SMP/E will write PTFs to one or more spill data sets.

SMPPTS spill data sets are permanent, partitioned data sets which you allocate and manage just like the SMPPTS data set. You then tell SMP/E about these spill data sets by defining them with either a DD statement or a DDDEF entry, using the reserved ddnames of SMPPTS1, SMPPTS2, SMPPTS3, and so on up to SMPPTS99 (99 spill data sets should be enough for everybody!).

Whether your SMPPTS spill data sets are defined by DDDEF entries or DD statements, SMP/E will refer to each of the data sets, in sequential order, when writing to, reading from, and deleting PTF members from the SMPPTS data sets. Also, after SMP/E writes a PTF member to the SMPPTS or a spill data set, SMP/E does not need to keep track of which data set that PTF member resides in. This means you are free to manage the SMPPTS and spill data sets however you want. You can merge data sets, split data sets, and move members from one data set to another, without affecting SMP/E.

We feel the SMPPTS spill data set function is so useful that we've

even made it available to current releases of SMP/E in PTF UR52518 for OS/390 Version 2 Releases 7, 8, 9, 10, and z/OS Release 1, and PTF UR52517 for OS/390 Version 2 Releases 5 and 6. Check it out now if you haven't already. By the way, for more details about the SMPPTS spill data sets, and other suggestions for better managing the size of your SMPPTS, check out the following Web site: **http://www.ibm.com/ servers/eserver/zseries/zos/ smpe/smppts.html**

HOLDDATA research

In a recent survey, when asked the question "Which aspect of preventive maintenance requires the most time and the most skill?," the most frequent response was "Researching System HOLDATA!" Not surprising. The first task in researching HOLDDATA is the simple act of finding the appropriate HOLDDATA text you need to read. SMP/E does a good job of telling you which HOLDDATA you should read, but it does not actually show you the HOLDDATA. You are forced to hunt for the actual text by browsing the SMPPTS, using the Query dialog, or even running the LIST command.

Therefore, to help you find the HOLDDATA, SMP/E now provides you with the HOLDDATA text you need to read by imbedding it directly in the SMP/E output. There are three additional reports produced by SMP/E during APPLY

and ACCEPT command processing to show you the HOLDDATA text:

1. The first report contains all of the **unresolved HOLDDATA.** Unresolved HOLDDATA causes a PTF to fail during APPLY or ACCEPT processing, and typically refers to ERROR HOLDDATA (PEs and Hipers).

2. The second report contains all of the **bypassed HOLDDATA.** Bypassed HOLDDATA is that which you deem OK by using the BYPASS operand on the APPLY or ACCEPT command, and typically refers to SYSTEM HOLDDATA.

3. The third and last report provides a **summary of all of the REASON-ids** from all of the HOLDDATA that appears in the first two reports.

As you might imagine, these reports can become extremely large, and, therefore, you may want to exclude some of the information from the SMP/E output. If that is the case, you can tell SMP/E to suppress certain kinds of HOLDDATA in the reports. You do this by specifying the REASON-ids for HOLDDATA to be suppressed with a new subentry in the OPTIONS entry called **Suppress Holddata** (SUPPHOLD). For example, you could tell SMP/E to suppress all DOC HOLDDATA in the reports. The resulting SMP/E output from an APPLY or ACCEPT command will provide you with all the HOLDDATA you need to read and understand.

HFS data sets...Which ones did I just update?

Are you sure you updated the correct Hierarchical File System (HFS) during your last APPLY? Depending on your cloning procedures and onto which directories you mount the HFS data sets, you could be updating the wrong file system. Well, SMP/E is here to help. SMP/E will now tell you exactly which HFS data sets are updated during processing. In the **File Allocation Report,** SMP/E already shows you the directories in the Hierarchical File System updated by SMP/E. Now, this report also shows the HFS data set where that directory resides. So, after you apply that important PTF, you can be sure you've mounted the right data set and updated the right file system.

Summary

In summary, SMP/E brings plenty of goodies to market in October...and not just for z/OS Release 2. There's help for that overflowing SMPPTS, easy reading for HOLDDATA, and explicit HFS data set identification just to name a few. There is also network packaging and delivery to complement SMP/E software installation—see the article *"SMP/E: It's not just for tapes anymore!"* for more. •

SMP/E: It's not just for tapes anymore!

KURT QUACKENBUSH

In October of this year, many enhancements will be available in SMP/E to simplify and extend the installation and service of software for z/OS and OS/390 systems. As the title states, SMP/E isn't just for tapes anymore. You'll now be able to receive input from a network location, in addition to traditional HOLDDATA, and associated materials such as documentation, samples, and text files (all the stuff you usually get on a tape). A companion for GIMZIP is the **GIMUNZIP** utility, which is used to recreate the data from a portable GIMZIP package.

information each data set contains. For example, does the data set contain MCS, HOLDDATA, or RELFILEs? This is important to know later during RECEIVE command processing so SMP/E can handle the information properly.
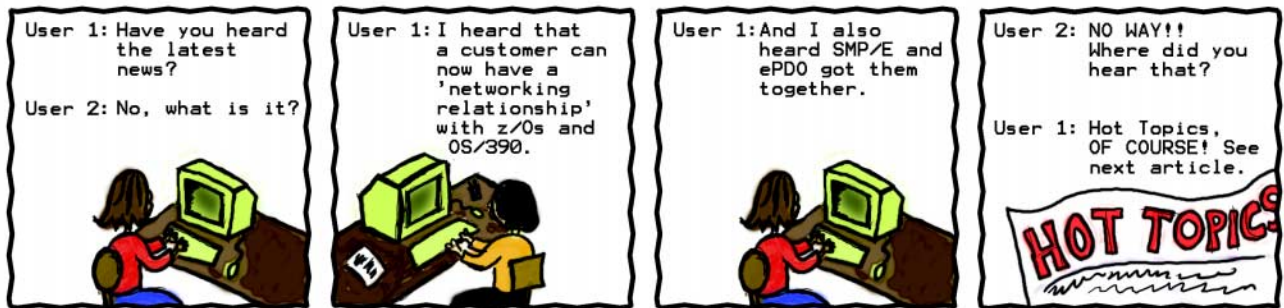
tape or DASD inputs. Also, there are other improvements in the area of SMPPTS management, HOLDDATA research, Hierarchical File System (HFS) management, and more (see the article *"Spreading the love: SMP/E for z/OS and OS/390"* for the rest of the scoop).

Receive from a network

You might think it's about time SMP/E entered the network world. Well, youre right. To more seamlessly integrate the electronic delivery of software with the SMP/E installation of that software, SMP/E will now receive input directly from a network location— either over the Internet, or even over your own enterprises intranet. This could be a great way to distribute software electronically throughout your enterprise.

There are two distinct components of this new capability in SMP/E. First is a packaging service utility. This utility, called **GIMZIP,** creates portable packages of software and associated materials. Typically, the packages will contain SYSMODs, RELFILE data sets,

The second component of this capability is the SMP/E network RECEIVE function. This is a variation of the SMP/E RECEIVE command, called FROMNETWORK. **RECEIVE FROMNETWORK** transfers a portable GIMZIP package across a network, extracts the information from the package, and then performs the traditional RECEIVE operations such as updating the SMPPTS and global zone and creating SMPTLIB data sets.

There are several steps to using this new SMP/E capability:
1. Create a GIMZIP package.
2. Receive the package.
3. Process any associated materials.

Create a GIMZIP package

A GIMZIP package would be created by IBM, or by anyone that wants to deliver software over a network. To create a GIMZIP package, you run the GIMZIP utility and tell it which data sets you want to include in the package. You can specify partitioned and sequential data sets, and you need to tell GIMZIP what kind of

Each of these data sets is transformed by GIMZIP into an **archive file**, which is the portable image of the original data. An archive file is a compressed image, and is therefore not directly viewable. However, you can also tell SMP/E that a data set contains text information that is to remain readable, called a **readme file.** A readme file is not transformed into an archive, but it does become part of the package.

In addition to the archive and readme files, GIMZIP also creates a packing list called the **package attribute file.** The package attribute file identifies the archive and readme files that compose the package, and contains an SHA-1 hash value for each of the files. The hash value is used for data integrity purposes and is checked later during RECEIVE command processing. GIMZIP (and RECEIVE) uses Integrated Cryptographic Services Facility (ICSF) to calculate hash values; therefore, youll need to configure and start ICSF. Refer to the ICSF System Programmers Guide for details.

zFavorites CD-ROM
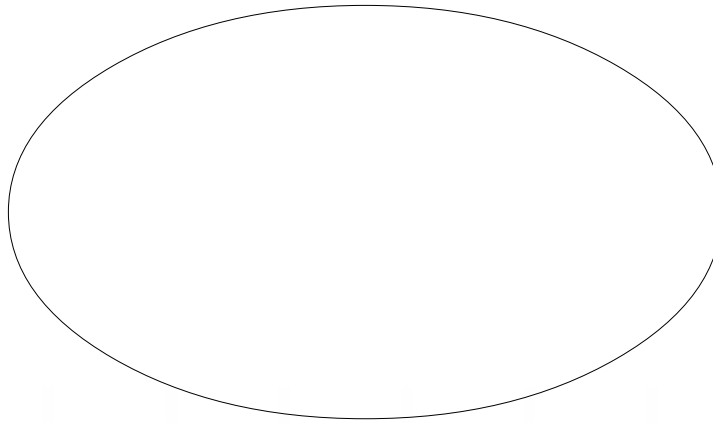
Diane Attebury

Hey! It's the zFavorites for zSeries credit card CD! You're gonna love this—it has all sorts of helpful Web links, like for:

Hardware
Operating Systems
Software
Language and Tools
ISV Development and Applications
Product Documentation
Marketing Information
Education
Support
Links to FREE downloads

Plus, it contains PDF versions of the z/OS and OS/390 Information Roadmaps and a preview of the Web sites linked to from the CD.

Also, don't forget—it has all existing issues of the Hot Topics Newsletter (including THIS issue) right on it for you to take anywhere!!

To use the CD:
Insert it in any standard CD-ROM and it should start automatically.

If it does not, then click on the **Start** button, choose **Run**...
and then type **x:\index.htm** (where x is your CD-ROM drive letter)
and press **Enter**.

Notes and Doodles

TECH JUMBLE

Unscramble these four word jumbles, one letter for each square, to form four ordinary words.

CEAEP

▢ ⊙ ▢ ⊙ ▢

VLEO

▢ ⊙ ▢ ▢

DAN

▢ ▢ ⊙

NXLUI

⊙ ▢ ▢ ▢ ▢

Answers on the back of this page

Then arrange the circled letters to form the answer to the question down below, as suggested by this cartoon.
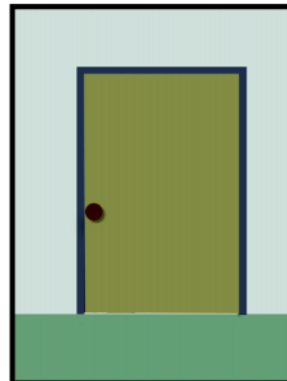
This operating system ain't what?

⊙ ⊙ ⊙ Ⓢ ⊙ ⊙

---

1. Please indicate which editions of the Hot Topics Newsletter you have read.
❑ OS/390 Release 8 (9/99)   ❑ OS/390 Release 9 (3/00)
❑ OS/390 Release 10 (9/00)   ❑ z/OS Release 1 (3/01)

2. What is your overall satisfaction with the Newletter?
❑ Very satisfied ❑ Satisfied ❑ Neutral ❑ Dissatisfied ❑ Very dissatisfied

3. How do you prefer to access and read the Newsletter? Please choose one of the following:
❑ Hardcopy ❑ Softcopy on CD or local repository
❑ Softcopy on the IBM Internet site or your intranet site
❑ All (depending on where I am and what I'm doing)

4. How useful would it be when searching the entire product library to have the Hot Topics Newsletter included in the search?
❑ Very useful ❑ Somewhat useful ❑ Neither ❑ Not very useful
❑ Not at all useful

5. The Newsletter is published two times a year. How often would you like it to be published? Please choose one of the following:
❑ Yearly   ❑ Two times a year   ❑ Four times a year   ❑ No opinion

6. How satisfied are you with the format and layout of the Newsletter (easy to navigate, pleasant to look at, etc.)?
❑ Very satisfied ❑ Satisfied ❑ Neutral ❑ Dissatisfied ❑ Very dissatisfied

7. How satisfied are you with the writing style of the Newsletter (enjoyable to read, keeps your interest)?
❑ Very satisfied ❑ Satisfied ❑ Neutral ❑ Dissatisfied ❑ Very dissatisfied

8. How do you obtain hardcopy issues of the Newsletter?
❑ With the product ❑ At SHARE ❑ At other conferences
❑ Purchase seperately ❑ Other (specify) _____

9. What one action should we take to improve your satisfaction with the Newsletter?

_____

_____

_____

Name _____
Company _____
Job Description _____
Phone _____
e-mail _____

May we contact you?
❑ Yes ❑ No

Thank you for your comments!

IBM

TECH JUMBLE

CEAEP

P E A C E

VLEO

L O V E

DAN

A N D

NXLUI

L I N U X
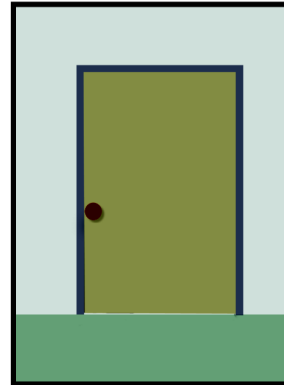
This operating system ain't what?

C L O S E D

||||
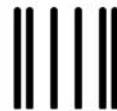
BUSINESS REPLY MAIL

FIRST-CLASS MAIL   PERMIT NO. 40   ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
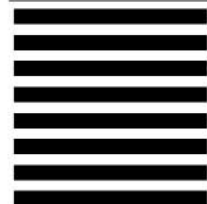
Department 58IA, Mail Station P384

2455 South Road

Poughkeepsie, NY 12601-5400

All of the files associated with a package—the archives, the readmes, and the package attribute file—are stored by GIMZIP in a directory in the hierarchical file system. This directory is referred to as the **package directory.** This package directory must then be made accessible to a network. SMP/E uses FTP to transport a package; therefore, the package directory must be made available on an FTP server. The FTP server could be the same z/OS or OS/390 system used to create the GIMZIP package, or it could be a different system, or even a different platform, where you stage the package directory. In any case, it is from this FTP server that the RECEIVE FROMNETWORK command will access and transfer the package.

After a package is staged onto an FTP server, the name of the server, and the location on the server where the package directory is stored, must be given to anyone that will be receiving the package.

Receive a package

After the package is created and staged to an FTP server, SMP/E can then be used to remotely get and process the package. This is where the RECEIVE FROMNETWORK command comes in. RECEIVE FROMNETWORK transfers a portable GIMZIP package through a network using FTP, stores the package on the local z/OS or OS/390 system, extracts the information from the package, and then optionally performs the traditional RECEIVE operations such as updating the SMPPTS and global zone, and creating SMPTLIB data sets.

You must tell RECEIVE FROMNETWORK what FTP server contains the package (server name or IP address), where the package is stored on the FTP server, and what the hash value is

for the packages attribute file. This server information must be supplied to you by the creator of the package. Depending on your local environment, you also may need to tell SMP/E how to navigate your local firewall.

During RECEIVE FROMNETWORK, SMP/E transfers the package files through the network from the FTP server to the local client and stores them in a directory within the local clients hierarchical file system. This new package storage directory is called the **SMPNTS,** which stands for SMP/E Network Temporary Store. This SMPNTS directory is similar to your tape library for physical media—just as a tape is staged from the loading dock to the tape library in preparation for RE-CEIVE processing, so is a GIMZIP package staged from the network to the SMPNTS directory. The SMPNTS is a simple directory that you identify to SMP/E using either a DD statement or DDDEF entry with the reserved ddname of SMPNTS. For example, this directory could be /u/smpe/smpnts/, or whatever suits your environment.

After all package files have been transferred, SMP/E can extract the data from the archive files and perform traditional RECEIVE operations in one step. However, this extra processing is optional, and you also can tell SMP/E to transfer the package and stage it to the SMPNTS, but not update the SMPPTS or global zone. You do this using the **TRANSFERONLY** operand on the RECEIVE FROMNETWORK command. The obvious companion to the TRANSFERONLY operand is then the ability to perform RECEIVE processing on a package that is already staged in the SMPNTS directory. This is done with another variation of the RECEIVE command called FROMNTS.

Whether you use RECEIVE FROMNETWORK to transfer and receive a package or RECEIVE FROMNTS to receive a package previously transferred, you have all the same options on the RECEIVE command that are available when you receive input from tape—select just SYSMODs or just HOLDDATA, filter by FMID, assign a source-id, and so on.

Process associated materials

When the RECEIVE command is completed, there may be associated materials in the GIMZIP package that are not consumed by SMP/E. Using a Custom Built Product Delivery Offering (CBPDO) as an example, there are archive files containing documentation from the DOCLIB and PGMDIR data sets, and there is an archive file containing sample jobs from the RIMLIB data set. After the SMP/E RECEIVE operation, these archive files remain in the SMPNTS and are in a compressed form. To extract the data from these archive files, you need to use the GIMUNZIP utility. GIMUNZIP extracts the data from an archive file and stores it in a new data set. You get to tell GIMUNZIP the archive files you want to extract, and, of course, you get to specify the names for the new data sets as well. As a bonus in CBPDO, a sample GIMUNZIP job to handle the DOCLIB, PGMDIR, and RIMLIB will be provided as a readme file in the package. All you need to do is update the sample job to specify your desired data set names and run it.

At this point, you would continue with the installation tasks

just as if the software was received from a physical tape.

Summary

In summary, SMP/E brings plenty of goodies to market in October...and not just for z/OS Release 2. Of course there is network packaging and delivery to complement SMP/E software installation. There is also help for that overflowing SMPPTS, easy reading for HOLDDATA, and explicit HFS data set identification—see the article *"Spreading the love: SMP/E for z/OS and OS/390"* for more. •

SQUEAK SAYS...
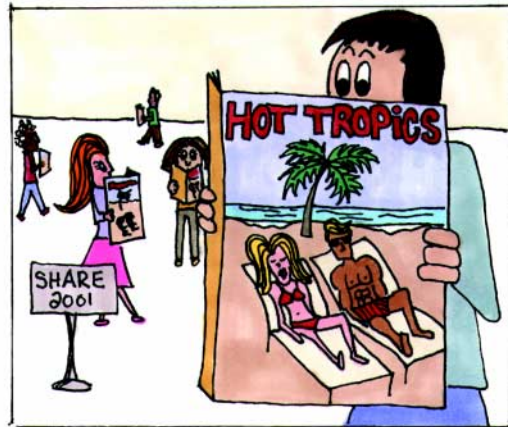
"Oops! I think we got the wrong shipment"

CBPDO in the fast lane

DEBBIE BEATRICE

And you thought you'd be handling those CBPDO product tapes until the end of time. Although you may be quite attached and comfortable with CBPDO tapes—you may even rely on them as a source of scratch tapes...change is afoot. CBPDO is moving into the realm of electronic

Liz Cicale

delivery to provide you with a means of getting your products faster by exploiting the new SMP/E function, **RECEIVE FROMNETWORK,** to leverage automation opportunities. With electronic CBPDO, or ePDO as we call it, you can order products and specify that you want to have them sent electronically. You can either download your order to z/OS or OS/390, or you can download it to a workstation.

ePDO's design has sought to include an end-to-end view of all the tasks...all the way from placing your order, tracking its status, downloading, installation, and redistribution.

Getting started

ePDO is available for both z/OS or OS/390 environments. ePDO uses SMP/E's **RECEIVE FROMNETWORK** function to download and optionally install an electronic package. SMP/E provides support for creating network-ready packages, manages the download of the package, and can **RECEIVE** the SMP/E installable

files. This SMP/E support is available in z/OS V1R2 and will also be made available to z/OS V1R1 and OS/390 license holders. Refer to the z/OS V1R2 announcement materials for availability information.

Deciding to use ePDO

ePDO is ideal for ordering products outside of major upgrades of your operating system or if you have forgotten to order a product and need it right away.

One of the first things you need to evaluate is the transfer rate capability for your connection to the Internet. The average ePDO package contains about 3-4 products, and is about 500 MB (compressed), if you use the enhanced CBPDO option to limit the amount of service delivered with your order. If you download directly to z/OS or OS/390, the sample jobs provided with your ePDO order can easily be run in a batch environment, eliminating the need to watch the bytes transfer.

Also, ePDO packages are comprised of a series of individual files, not a single large file. This helps to address situations where the download may have been interrupted. Transfer can automatically restart at a file level, not from the very beginning of the download operation.

Ordering an ePDO

ePDO orders for either z/OS or OS/390 platform products must be placed using ShopzSeries:
https://inetsd01.boulder.ibm.com/ ShopzSeries/ShopzSeries.jsp

What products can you order? Any of the products supported in CBPDO...with two exceptions— neither z/OS nor OS/390 base or optional features are available for ePDO orders. Why? ServerPac is the recommended delivery option for these products because it provides a substantial installation setup. We're looking at electronic delivery of z/OS or OS/390, but we have more work to do.

You'll be able to track your order's status from ShopzSeries. When your order is ready, you'll receive an email that contains information about your package and a link to a Web page customized to your ePDO order.

Your ePDO order is ready

For a time, your ePDO package will consist of both an electronic and a physical delivery. The electronic package is a combination of information available on a Web page specific to your order and the electronic ePDO package itself. Your ePDO download package contains all the contents available today on a CBPDO tape, with samples and documentation updated to reflect electronic delivery. The ePDO order will remain on the IBM server for 30 days.

When your ePDO order is available, you access your order-specific Web page through the ShopzSeries Web page. A customized Web page will contain:
- Installation instructions
- A packing list
- Sample SMP/E **RECEIVE FROMNETWORK** job, customized to your order
- A launch point if you choose to download the ePDO package to a workstation.

The physical components of your order will be delivered as they are today and will include a packing list, CDs for both workstation client code and publications, entitled hardcopy, and informational memos. We are working with other groups within IBM to enable these physical components for electronic delivery.

Downloading your order

You have two download options. You can either download directly to a z/OS or OS/390 system, or you can download to a workstation as an intermediate node. For either method, you can use the sample **RECEIVE FROMNETWORK** job provided on the order's Web page to help you with the installation.

Choosing a download method

Depending on your enterprise security policies, network configuration, or preference, you need to decide whether you will download your ePDO package directly to your z/OS or OS/390 system (using FTP as the transfer protocol and userid and password authentication), or if you will download the package to a workstation as an intermediate node (using an browser-based, HTTP-initiated, signed Java applet to download). Both methods use a hashing algorithm to provide data integrity.

Requirements for downloading directly to z/OS or OS/390

In addition to the new SMP/E, there are two basic prerequisites for ePDO if you are downloading to z/OS or OS/390:
- Configuration of Integrated Cryptographic Services Facility (ICSF)
- Identification of firewall navigation information.

ICSF Requirements

If you are downloading directly to z/OS or OS/390, or if you are forwarding using SMP/E **RECEIVE FROMNETWORK** to access files from a workstation, you'll need to configure Integrated Cryptographic Services Facility (ICSF). For those of you who may not already be using ICSF, it's an element of both z/OS and OS/390. ICSF works with the hardware cryptographic feature and Security Server RACF element to provide secure, high-speed cryptographic services. SMP/E uses ICSF services as a way of providing data integrity of your ePDO package. ICSF has both hardware and software setup requirements.

The hardware requirement is for the Cryptographic Coprocessor feature.

The software customization requirements are to:
- Order and install the cryptographic enablement diskette.
- Define at least one cryptographic coprocessor to each LPAR.
- Complete the remaining ICSF setup.

Refer to *ICSF System Programmer's Guide, SC23-3974*, for more information about the specific hardware features and software setup.

Firewall requirements

You will also need to identify information to help you navigate your firewall to download your ePDO order. This information will be used by the SMP/E to initiate the transfer.

Requirements for downloading to a workstation

If you download your ePDO order to a workstation as an intermediate node, you'll need:
- A Java-enabled browser with the ability to accept a signed Java applet.
- The ability to share or transfer data to z/OS or OS/390.
- Sufficient storage available for the compressed package.
- Firewall navigation information.

You can then use the method best-suited to your environment to either transfer the files (you may choose SMP/E **RECEIVE FROMNETWORK** or FTP), or to access the files from z/OS or OS/390. For example, you could use a network file system (NFS) or journaling file system.

If you use **RECEIVE FROMNETWORK,** which uses the **GIMUNZIP HASH=YES** parameter, the z/OS and OS/390 requirements are the same as those described for downloading directly to the host. This method also requires that the workstation be configured as an FTP server.

Downloading to a z/OS or OS/390 node

First, update the sample SMP/E **RECEIVE FROMNETWORK** job provided on the Web page for your order. This sample job processes the SMP/E installable files and is customized to reflect information about your package; however, you need to update the job to reflect your network definitions and installation preferences. Based on your updates to the sample job, you can either:
- Download and **RECEIVE** your order.
- Download your order using the **TRANSFERONLY** option, delaying the **RECEIVE**.
  If you choose the **TRANSFERONLY** option, you then use SMP/E **RECEIVE FROMNTS,** which will uncompress the SMP/E files and do the **RECEIVE**.

The ePDO package itself provides another sample job to

process non-SMP/E files with your package that remained compressed. Some of these files are RIMLIB and DOCLIB. The README file will point you in the right direction to continue the install. The README file is in EBCDIC format and viewable on the host.

After the CBPDO tasks are done, you complete the product installation using the instructions in the program directories.

Downloading to a workstation as an intermediate node

This method uses a browser-based, HTTP-initiated transfer. Your package is stored on the local system ("intermediate" system) where the browser is running.

First, select the "Download to workstation" link on your order page to launch the download, which launches an applet. After the download is complete, you must make the package files accessible to SMP/E to initiate the installation on a z/OS or OS/390 system. There are two methods which can be used to make the files accessible to SMP/E:

* Use SMP/E **RECEIVE FROMNET** to load the files directly from your intermediate system, using FTP. For this, your workstation must be configured as an FTP server.

You can also use the sample **RECEIVE** job supplied on the order's Web page to assist with this step.

* Use SMP/E **RECEIVE FROMNTS.**

And in conclusion...

When you use ShopzSeries and ePDO, you can benefit by having a more seamless integration of the ordering and installation process. You may also find that you are able to order products closer to the time that you actually want to deploy them in your environment.

The ePDO pilot is scheduled to begin 4Q/2001. We'll be communicating more details about the pilot in the coming months. If you have questions or would like more information, you can contact me, Debbie Beatrice, at *debbiebe@us.ibm.com* •

Shop zSeries: Order your zSeries and S/390 software over the internet

STEPHEN SAROKA

https://inetsd01.boulder.ibm.com/ShopzSeries/ShopzSeries.jsp

Highlights:
* Tired of sorting throught product requisites, checklists, etc., trying to figure out the product migration paths.
* ShopzSeries (formerly ShopS390) provides a new, self-service interface to IBM for planning and ordering zSeries software faster and easier over the Internet.
* ShopzSeries automatically analyzes your SMP/E CSI installed inventory.
* Provides you a preselected IBM software product catalog list of your upgrade options.
* Performs technical requisite checking.
* Submits your order online and tracks the order status through delivery.

Kerry Collia

What is ShopzSeries?

Planning and ordering software upgrades used to be a time consuming and complex task. Just trying to figure out what to order to upgrade your current system required lots of research.

Now, there is a fast, easier way to select the desired upgrade offering such as a ServerPac or

CBPDO and order your z/OS and OS/390 software (including any of the subsystems like DB/2, IMS, CICS, etc...). Just visit the ShopzSeries Web site and upload your SMP/E CSI installed inventory from your target system. ShopzSeries automatically analyzes your installed inventory and maps it to your IBM licensed inventory records and the latest IBM software product catalog. It gives you a preselection of your upgrade options including migration from OS/390 to z/OS!

Then if desired, it lets you
* Add additional products to the order list
* Remove products from the order list

- Perform technical requisite checking and various "what ifs."
- Save the order and complete later or if done, submit the order to IBM.

After submitting your order you can track the order status through delivery and your order is saved on ShopzSeries for your future reference.

Additional capability within ShopzSeries includes the ability to view your IBM licensed inventory (with or without your installed inventory and the latest available products) and you can also view version upgrade opportunities across your entire establishment.

What is the customer reaction to ShopzSeries to date?

Since its initial deployment in late summer of 2000, this new self-service capability has been a big success with customers. Both customer satisfaction and usage have been excellent.

In just a few months, this new solution has already captured 40% of the total US ServerPac volume.

Customer satisfaction is also high, running about 90% satisfied!

How can you get started using ShopzSeries?

Customers can access ShopzSeries Web site directly at **https:// inetsd01.boulder.ibm.com/ ShopzSeries/ShopzSeries.jsp** to learn more about this new and exciting e-business solution!

First time users need to register at the ShopzSeries Web site (or if you like you can try out ShopzSeries in "guest mode" although you will not be allowed to actually submit any orders until you are registered and have been authenticated)

How does ShopzSeries work?

After completing the registration at the ShopzSeries Web site, you will need to generate a SMP/E Report of your "installed" system inventory needing upgrades, and make it available to ShopzSeries. ShopzSeries provides detailed instructions and an easy way to do this.

You then select the desired upgrade offering (CBPDO or ServerPac) and desired "view" of the IBM z/OS and OS/390 software catalog.

ShopzSeries presents an analysis of your "installed" inventory against the current IBM software catalog and your licensed inventory.

Upgrades are identified *(including new versions like z/OS!)* for you. In the case of ServerPac, an entire replacement order is preselected.

You can also "shop" for additional products and add these to your order.

You can perform technical requisite checking and "what ifs". Once you're satisfied with your order content and it meets technical requirements, you submit it.

You will receive an e-mail "return receipt" indicating your order has been received by IBM, and you can track order status through delivery.
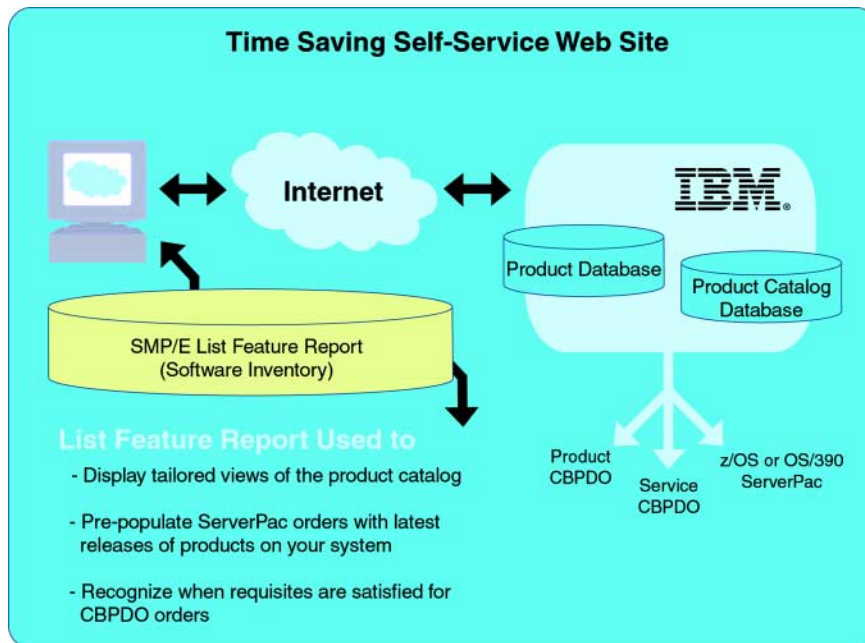
Is ShopzSeries available in my country?

Currently, ShopzSeries is generally available in the US and much of Western Europe. ShopzSeries will also be made available to other countries this year. Please stay tuned to the ShopzSeries web site at **https:// inetsd01.boulder.ibm.com/ ShopzSeries/ShopzSeries.jsp** •

Time Saving Self-Service Web Site

Internet

IBM.

Product Database

Product Catalog Database

SMP/E List Feature Report (Software Inventory)

List Feature Report Used to

- Display tailored views of the product catalog
- Pre-populate ServerPac orders with latest releases of products on your system
- Recognize when requisites are satisfied for CBPDO orders

Product CBPDO

Service CBPDO

z/OS or OS/390 ServerPac

Everything you always wanted to know about softcopy repositories
(and did not hesitate to ask)

JIM STEIPP

Books Books Books...

We have been inundating you with softcopy books. As we have responded to your need for more and more online books, the sheer volume has made it more difficult for you to access them. There are several thousand books for almost two hundred IBM server and mainframe products, and that doesn't include "red" books or any of the books you receive from other computer and software manufacturers.

This article will examine the various ways that you can manage your softcopy books. We will explore how you can develop an overall strategy that provides your users with the information they need in order to be productive. We will, however, focus on central repositories and try to help you understand some of the decisions you must make in preparation for setting up your own delivery system for softcopy information.

What are my alternatives for providing my users with softcopy information?

You can offer your users access to softcopy information in any of several ways. These include:
- The IBM Online Library collections on CD-ROM.
- The IBM Internet Library on the World Wide Web.
- Your own repositories of softcopy books.

What about CD-ROMs and the Internet?

CD-ROMs and the Internet are useful tools that are available to help you distribute and manage the softcopy information. They remain very popular methods for accessing IBM books. For many of you, they might be all you need. You and your users can access the IBM Internet Library at this URL: **http://www-1.ibm.com/servers/s390/os390/bkserv/**

Kerry Collia

You will need only a browser to access the entire z/OS, OS/390, z/VM, VM, and VSE/ESA™ and the "red books."

Remember, in certain circumstances, particularly when users must be mobile or must work in areas where they have no Internet access, CD-ROMs might be the only way in which they can view softcopy information.

Why would I want my own repositories?

There are three basic reasons for creating and maintaining your repositories:
- **It is not practical to use the CD-ROMs in the IBM Online Library collections.**
  In a large organization, you can't supply everyone who needs them with their own copy of all the CD-ROM collections and keep them current.

  Your network servers can't handle CD-ROMs or they can't accommodate the number of CD-ROMs that your users would require.

- **Your users do not have access to the IBM Internet Library.**
  For reasons of security or economy, you do not provide your users with Internet access. The performance of the Web (availability or speed) is insufficient to meet you users needs.

- **You want to maintain central control over the acquisition and distribution of softcopy information in your organization.**
  You want to be sure that all of your users have same correct, current information.

  You want control over the cost of acquiring and distributing information.

  (Neither of these might happen if everyone is responsible for obtaining and distributing books on their own.)

How many repositories can I have?

You need only one repository to enable the SoftCopy Librarian to manage any of your bookshelves and books. However, you can set up as many repositories as you would like, especially if any of the following situations apply to you:
- All your users do not have access to the same host or server.
- Users on the same host might not have access to the same LPAR.
- You have a large number of books that are easier to manage in unique sets (for example,

IBM and non-IBM).

- When migrating from one version or release of a system to another, you keep the later books in separate repositories (to make it easier to delete the older sets).
- When you run a test release of your system along with your production release, you keep the books for each release separate from each other.

When deciding on the number of repositories to set up, also keep in mind the following:
- The number of books to which you want access.
- The needs and interests of your users.
- Your need to control access to certain libraries.

What does the SoftCopy Librarian manage?

The SoftCopy Librarian actively manages the following types of files:
- BookManager Books (BOOK or .BOO files) Bookshelves (BKSHELF or .BKS files) Bookshelf indexes (BKINDEX or .BKI files)

Active management is more that just transferring files to your repositories, it also includes:
- Maintaining a catalog of the contents for each of your repositories.
- Comparing what's currently in your repositories with new sources of bookshelves and books.
- Providing you with the comparative status of the items in a source versus a repository.
- Updating the catalogs whenever you transfer new bookshelves and books to your repositories.
- Identifying items in your repositories that can be removed because there are newer versions of these items also in your repositories.

The SoftCopy Librarian can transfer any type of file to any repository you choose but it cannot provide active management for the files. You must manage and access these files yourself.

Where can I locate my repositories and how can I access the books?

There are several ways in which you can locate and access your repositories, depending upon how you set them up.

On a z/OS or OS/390 host, you can access the books with:

In sequential data sets:
- BookManager READ/MVS from an attached 3270-type terminal or through a 3270-like emulator session from an attached workstation.
- BookManager BookServer for MVS from an attached workstation through a Web browser

In hierarchical files:
- BookServer from a workstation communicating over the Internet or your intranet through a Web browser.

On a Linux, AIX®, or other Unix-type server, you can access the books with:
- BookServer from a workstation communicating over the Internet or your intranet through a Web browser.

On a LAN server, you can access the books with:
- Any BookManager Library Reader from an attached workstation.

On a local PC workstation, you can access the books with:
- Any BookManager Library Reader™ from an attached workstation.

On your own personal computer, you can access the books with:
- Any BookManager Library Reader from your workstation.

How can I arrange my repositories?

There are several ways in which you can arrange your repositories:
- **Put all your books into one central repository.**
  This is probably the simplest way to maintain your library. You have only one place to put books, regardless of from where they came. You don't have to worry about updating more than one repository or coordinating between them. It also avoids maintaining duplicate books. However, this will be slower for the SoftCopy Librarian, and it might be more difficult for your users to find specific books. In addition, all the books might not all fit onto one DASD device. The SoftCopy Librarian allows you to specify the device and data set or directory to use to install new books, but it will not do any load balancing, nor will it automatically put an updated book onto the same device as the book it replaces.
- **Clone one repository onto several platforms.**
  If all your users cannot access one central repository, you can create identical versions of a repository and place the copies on the different platforms that are accessible. While this solves the access problem, it means that you now have several repositories to maintain simultaneously and you will have duplicate copies of all the books on each platform. The SoftCopy Librarian can help you maintain several identical repositories through its ability to create "clones" of one repository that you maintain as a "model" for all the others. However, you may be faced with space problems on one or

more of the platforms, and performance may still be slower than you might want.

Note: You can also "synchronize" or merge the contents of several separate repositories. This will ensure that any books that appear in any repository will also appear in all the others. This is useful if you have more than one person maintaining repositories independently from each other.

- **Put unique sets of books into separate repositories.**
  This would be faster than putting everything into one or more of the repositories and it would save some space. You could put all IBM "white" books into one repository and "red" books into another. Or, you could put all IBM books into one repository and your non-IBM books into another. This might be the best balance between the two previous arrangements. There probably

won't be any duplication of books, and each repository will be smaller. If there are any IBM collections that don't share books, (such as the S/390® red books), those collections could also be separate. To do this, you could use directories such as */u/softcopy/ibm, /u/softcopy/red,* or */u/softcopy/oem* for those repositories.

Note: The alternative to putting each CD in a collection into a separate repository is not practical because of the extreme duplication of books and because bookshelves move from disk to disk between editions of a CD collection.

What information must I gather before I can define a repository?

You must decide upon the arrangement of your repositories:
- How many do you plan to have?
- What will they contain?
- Where will they be located?

Then you can provide the following information when you define each repository:

For all types of repositories:
• A name for each repository.
  Note: This name can be anything you want. It is the name that the SoftCopy Librarian uses on your workstation to identify your repositories.
  • One or more directories or high-level qualifiers that will contain the book (.BOO) files.
  • A directory or high-level qualifier (HLQ) that will contain all shelves. It can be the same directory or HLQ that you specified for the books and indexes.

- A directory or HLQ that will contain index (.BKI) files. It can be one of the directories or high-level qualifiers that you specified for the books.
  Note: For host repositories in sequential data sets only: Your bookshelves and books do not require unique HLQs. For example, if your BookManager code is installed under IBMBK, the bookshelves and books could be installed under IBMBK.SOFTCOPY.

For host repositories in sequential data sets only:
- A HLQ for the bookshelf list (EBRMGMT).
- The type of DASD device you are using.
- The block size for that DASD device (provided).
- The volume in which the data sets will reside.
- The applicable data, management, and storage classes that you are using *(for SMS users only)*.

For host, Linux, AIX, or other Unix-like repositories only:
- The name or IP address and the user ID and password you use to access your host or server.

For host, Linux, AIX, other Unix-type, and LAN server repositories only:
- Any FTP commands required by your location to transfer the bookshelves, books, and indexes from the workstation to your host or server.

Are you ready to take the next step? Yes!

Let the SoftCopy Librarian manage your repositories for you. For more information about system requirements, obtaining the SoftCopy Librarian code, and installing and using the SoftCopy Librarian, refer to the SoftCopy Librarian User's Guide on your CD collection or in the IBM Internet Library. •

Are you searching for a better way to look up z/OS, OS/390, and now VSE message explanations?

Then step right up ladies and gentlemen! Don't crowd! Have I got an deal for you!

You might have already seen the LookAt message help that is currently available on the Web. It's the help facility that doesn't need help to help you. You are not expected to know anything about the library or how to construct a search argument. You only need to know the number of the message that you want to find and select a release.

Well, I'm prepared to go one better.

Today and today only, to readers of HotTopics, we are able to let you in on this \ special offering. Starting June 30, IBM will make available, on the IBM Web site, new, improved version of LookAt. You can get there using the zFavorites for zSeries mini-disk that's included in this issue.

Just put in your CD reader,
        click on the "LookAt..." Tab,
                and bookmark the link.

When you get there, take a look at LookAt.
I think you'll like this one.

We're making LookAt better...by listening to you

JIM STEIPP

You asked for a better way to retrieve information from our z/OS and OS/390 library.

We gave you LookAt for messages.

We know that you are using it, because at the beginning of the second quarter, it made the top 10 list of most-visited IBM Web sites.

You have asked us, through our feedback page and surveys we sent you, to make some more improvements in the way LookAt works and to expand the number of messages that are available.

Well, we have expanded the number of messages that are now available. We've already added VSE message explanations, and we're working on VM messages.

We are also about to roll out a new look for LookAt on the Web. We are trying to make the user interface clearer and easier to use. To do that, we actually made two prototype versions available to a selected set of real customers, like you. We asked the customers to try out the prototypes and let us know what they liked and to make recommendations for improvements. We selected our testers from a list of customers who had sent us

Paul Rowntree

feedback and who participated in an extensive survey that we did at the last SHARE. The new interface will look something like the screen capture included with this article.

This is very close to the final version. We are still fine-tuning it based on the suggestions we've received. It does give you a good idea of how it's organized and how it will look. We have tried to arrange things more logically on the page and also differentiate more clearly both libraries that have been enabled for LookAt and older libraries that continue to rely on traditional BookManager searches. You will see new libraries as we continue to expand the current offerings as additional messages books are enabled. This new LookAt will be available concurrently with z/OS V1R2.0 by the end of June. Please try it out when it is available, and let us know what you think by using the feedback button at the bottom of the LookAt page.

You also asked us to provide LookAt access to return codes and

some other things like commands, macros, language statements, and few more. We are working on these and also a version of LookAt that will run on your workstation or laptop and use CDs in addition to the Web so you will be able to take

LookAt with you, wherever you go. These are still in the future.

However, by the end of the third quarter, we do plan to have an "all-books" choice that will let you list all the books currently enabled for LookAt and, for ISPF only, a "point-and-shoot" capability so you don't need to type the message IDs. We are looking at the feasibility of extending this capability to the Web site, but that will probably take us a while longer.

In the meantime, thanks for using LookAt. We will continue to improve it, so you can continue to look at the information you need.

Note: If you haven't already bookmarked the LookAt page, use the included zFavorites CD to find your way there. •

New IMS enhancements for OS/390 and z/OS

BARBARA KLEIN

There have been a number of new IMS enhancements since the October 27, 2000, General Availability of IMS V7 which continue to enhance the capability of OS/390 and z/OS. These are being provided through the service process and/or with new/enhanced IMS Tools.

IMS enhancements being provided through the IMS service process for
Program number: 5655-B01
1. 64-bit support

- IMS V5, V6 and V7 were all enabled to run on the new z/OS systems. This support was provided with the GA tape of IMS V7 and provided in the service stream, enabling APARs were available October 2000 for IMS V5 and V6.
- IMS V7 also provides page fixing of OSAM and Log buffers above the 2GB bar through APAR PQ42127 in April 2001. Expanded storage is now combined with central storage, thereby eliminating paging between expanded and central storage. All buffers for OSAM databases are page-fixed above the 2 GB bar. Use of OSAM buffers above the 2 GB limit is transparent to the users. If IMS is going to run in this environment, log buffer storage will only be fixed above 2 GB if the block size is a multiple of 4096. The virtual storage for these buffers remains below the 2 GB bar or in 31-bit mode.

2. Additional sense codes and information on IMS V7 messages for enhanced support of the zOS and OS/390 Communications Manager in January 2001

- IMS is sending a sign-on screen after receiving a LUSTAT X'082B' from a System Logical Unit 2 (SLU2) TELNET Extended Terminal Option (ETO) Terminal in APAR 42530. It addresses the customer requirement REQ00074249 where IMS should send a sign-on screen after receiving a LUSTAT X'082B' from a SLU2 TELNET ETO Terminal. For dynamic SLU2 TELNET terminals, when a VTAM LUSTAT X'082B' is received, IMS will Sign-off the terminal instead of sending a RESETSR. Module DFSCVGF0 has been modified to check for LUSTAT x'082B' for SLU2 dynamic (ETO) TELNET terminals. When an LUSTAT of x'082B' is verified, the node will be signed off and the connect message, MSGDFS3649A, will be issued. OS/390 V2.10 is required and IMS V7 APAR PQ41143 has to be installed.
- IMS is including the TCP/IP address in the IMS message in IMS V7 APAR PQ41143. The TCP/IP address and port number for Telnet terminal have now been included in the IMS Messages. This is only important for messages issued before a session has been successfully established. After a session has been established, the IP address can be extracted via the VTAM D NET,ID=nodename command. The TCP/IP address and the port number have been included in the DFS3672I message. Note: OS/390 V2.10 is required to extract the IP address and the port number. Module DFSCNXA0 has been modified to extract the IP address and port number at logon time from the CV64 (VTAM® control vector 64) and stores the information in the Communication Line Block (CLB). DFSCNXA will also include the IP address and the port number in the DFS3672I

message when available. Module DFSSIST0 has been modified to include the IP address and the port number in the DFS3650I message and to pass the IP address and the port number to the DFSGMSG0 exit.

3. IMS Support for DB2 MVS groups for Sysplex restart in January 2001 in IMS V7 APAR PQ42180

- IMS support for DB2 MVS groups is an IMS Dependent Region Group attach, providing Dependent Region access to available DB2, for Sysplex restart. IMS uses MVS™ Name Token Services to translate the DB2 "group attach name" into a DB2 subsystem name without requiring the DB2 V7 enhancements for group attach. IMS support of DB2 MVS groups, is not to be confused with DB2's Group Attach. They are separate entities. DB2 Group Attach does not rely upon IMS for any reason or purpose. IMS support of DB2 MVS groups only relies upon DB2 to the extent that, if the user does not specify a DB2 MVS group, IMS cannot select a DB2 from an MVS group. If IMS is unable to convert a user specified MVS group name into a valid DB2 name, an existing informational message will be issued to the IMS user.

4. Additional IMS V7 High Availability Large Database (HALDB) Performance and management enhancements. IMS Version 7 had enhanced, through independent partitioning to virtually unlimited, the capacity of the IMS full function databases. HALDB also provided enhanced availability of that data for those customers needing a 24X7 environment. Since the IMS V7 GA

additional enhancements have also been provided to improve performance and management of these databases:

- PQ37015 provides improvements in Secondary index migration performance
- PQ36991 provides improvements in Indirect List Data Sets (ILDS) creation performance
- PQ37127/UQ47735 provides improvements in DFSMAID0 performance
- PQ38822/UQ45056 provides RECON partition LIST command support
- PQ35893 adds the ability to initialize through Batch Commands the new HALDB databases.

Database Recovery Control (DBRC) Init command support is provided to initialize a HALDB and the associated partitions. The INIT.DB command is changed to add two new parameters: TYPHALDB and PARTSEL, and a new command, INIT.PART, is used to initialize the partitions.

5. JDBC access to IMS DB data from CICS TS/390 2.1 and DB2/390 V8.

- IMS V7 provided Java application support for running under the IMS Transaction Manager and accessing data through JDBC from IMS DB and/or DB2 databases. Enhancements are being provided for IMS Java support for enabling that JDBC access to IMS DB data to be extended to Java applications running in CICS Transaction Server V2R1 for OS/390 and z/OS or from Java Stored Procedures running under DB2 V7 for OS/390 and z/OS.

IMS Tools enhancements:

IBM IMS Command Control Facility for z/OS
Program number: 5655-F40
GA 1/01

- IMS Command Control Facility is a batch processor that can run as an IMS Batch Message Partition (BMP), IMS DL/I batch or standard MVS batch job. It can execute commands across as many as 64 IMS regions on any number of MVS images.

IBM IMS High Performance System Generation Tools for z/OS.
Program number: 5655-F43.
GA 1/01

- This High Performance System Generation (Sysgen) package includes three utilities. Fast Sysgen reduces time required to perform IMS system generation for application resource changes. Merge Clone ensures that each application resource is defined with compatible values across each copy of IMS. Sysgen Compare determines whether differences exist between two data sets.

IBM IMS Message Format Services Reversal Utilities for z/OS.
Program number: 5655-F45
GA 1/01

- This Message Format Services (MFS) tool helps improve database management and system performance through MFS Reversal, which recovers source and compare deltas, and MFS Compare, which helps you validate which MFS source library is operating in a particular IMS environment.

IBM IMS Image Copy Extensions for OS/390.
Program number: 5655-E10
GA 4/01

- IMS Image Copy Extensions obtains an image copy of the database and recovers a database data set. The tool also includes a space-saving compression feature. New enhancements include a dynamic allocation capability and a complete, one-step database HASH Check function.

IBM IMS High Performance Pointer Checker for OS/390.
Program number: 5655-E09
GA 4/01

- Used regularly, this tool helps ensure that your database pointers are error-free and alerts you when database reorganization is needed. This tool also assists in the analysis of a corrupt database, reducing diagnostic and repair time.

IBM IMS Online Recovery Service for z/OS V1R1.
Program number: 5655-E50
GA 4/01

- In a single pass of input data, this tool manages the accumulation and merging of recovery-related database data for the recovery of one or more IMS full-function databases or Fast Path areas. Data sets are processed in parallel, while multiple database data sets are updated simultaneously. A timestamp recovery function provides recovery to any prior point in time.

IBM IMS Database Control Suite for z/OS, Version 2.1.
Program number: 5655-F76
GA 6/01

- This tool is a day-to-day IMS database environment management tool for DBAs, system programmers and application support specialists. It enables them to collect and build on the data required to generate job control language (JCL) for database backup/recovery/reorganization management using IMS Base Utilities and IBM High Performance Tools.

IMS Fast Path Online tools V2R1.
Program number: 5655-E3
GA 7/27/01.

- This tool provides a new utility called DP Online Area Extender which allows you to

dynamically expand the size of a DEDB area online via expanding the SDEP portion or the IOVF portion of DEDB area, the benefit of this allows 24/7 availability of an IMS Fast Path DEDB area.

IMS Connect V1R1 enhancements provided through the service process for
Program number: 5655-E51
1. Local support, available via APAR PQ45057 4/01, for communication using Program Calls without requiring TCP/IP from a webserving application to IMS in a z/OS or OS/390 environment, easing the management in this environment. This support is used by the VisualAge® for Java (VAJava) IMS Connector for Java for creating Java applications that can access IMS transactions from WebSphere Application Servers for zOS and OS/390.
2. Unicode support, available via APAR PQ47906 5/01, for sending Unicode application data to an IMS host application capable of dealing with Unicode, such as a Java application running in IMS.
3. ACK/NAK required notification support, available via APAR PQ46195 4/01, provides client notification that an ACK or NAK response is required by the client without additional testing of data received. This notification will be sent to the Client in the CSM and RSM.
4. Output message structure change, available via APAR PQ48182 5/01, to include the full message length preceding the output message to the client, reducing the design and coding effort of a client application.

IMS Connect V1R2.
Program number: 5655 E51
GA 2H01
• WebSphere Adapter support, through the addition to IMS Connect in V1R2 of the IMS

Connector for Java J2EE Runtime support. J2EE is the industry-wide standard version of our IBM Common Connector Framework. This IMS Connector for Java J2EE support provides the runtime environment for Java applications developed using the VAJava IMS Connector for Java J2EE Development support. VAJava could be used to customize connectors and generate Java Beans that are deployed into the WebSphere Application Server or any other application server that conforms to the J2EE Connector Architecture. This IMS Connector support is designed to be used with the Java 2 Platform, Enterprise Edition, Connector Architecture Specification V1.0. IBM plans to make this IMS Connector for Java J2EE support, along with the other Websphere Adapters, available in the second half of 2001. General Availability depends on completion of the

J2EE Connector Architecture Final Specification and the J2EE Platform V1.3 Specification by JavaSoft. This specification is similar to IBM's proprietary Common Connector Framework (CCF) which was delivered in VAJava V2.0 along with the IMS Connector for Java for deployment in Websphere Application Servers. The specification of the Connector Architecture is currently in Public Draft Status. To download the draft specification, visit
http://www.javasoft.com/j2ee/ connector

Find out more about these and many other new and enhanced IMS capabilities for z/OS and OS/390 at our newly improved Web site at
http://www.ibm.com/ims •

What happened to ServerPac for z/OS?

JOHN EELLS

What did we do to ServerPac? Well, we've made more than a few changes. Now you can automate assigning data sets to volumes to get to IBM's recommended system layout in about 60 seconds. A new View and change facility gives you 28 built-in groupings you can use for all sorts of things. You can rename all the data sets if you want to. You don't have to use "logical volumes" any more. And initial setup is done for Basic Sysplex, System Logger, WLM Goal mode, SMS, TCP/IP, and OCSF, as well as preparation for IBM's new License Manager. All available since March, 2001.

A new look for Modify System Layout:

The first difference you'll notice about ServerPac is that the **Modify System Layout** dialog won't yell at you before you have a chance to change things!

Instead, you'll see this new **Modify System Layout options** panel. If it's been a year since your last installation you won't have to remember what a "SUMD U" command means. You can either select from options written in English or use the commands, which remain available.

```
CPPP605T ------------- Modify System Layout ( AM000321 ) ---------------------
OPTION ==> _

    A   Create a Recommended System Layout (Automatically assign target and
        DLIB data sets to physical volumes by data set type)

    C   View and change data sets by selected attributes
    T   View and change device type table (DEVT)

    D   Data Set Summary (SUMD)
    M   Merged Data Set Summary (SUMD M)
    S   Shipped and Merged Data Set Summary (SUMD S)
    U   User Data Set Summary (SUMD U)

    V   Physical Volume Summary (SUMP)
    L   Logical Volume Summary (SUML)

    P   Product, Feature and Element Summary

    ------------------- Session Control Options -----------------------
    K   Keep Changes made in this dialog session so far (SAVE)
    B   Back Out changes from this dialog session (CANcel)
    R   Reset all Values to Shipped (SHIP)
```

The old panel is still there, as Option P, but we think you will be more interested in using the two new options at the top.

The first one, *Create a Recommended System Layout*, will assign target and DLIB data sets to volumes automatically. The second, *View and change data sets by selected attributes*, is a powerful new facility for making changes to data sets.

View and change — slicing and dicing the Data Set Lists

Let's take a look at the new *View and change* option first. Before, you could group data sets only by assigning them to logical volumes. Then, you could assign the logical volumes to physical volumes. Now, the you can create a display using any of 28 attributes so changes to things like volume assignment can be done much more quickly. You select values within each attribute to get exactly the list you want, and you can save the each list in a file. Here is the list of attributes:

- Whether the data set is to be allocated
- BLKSIZE
- Current data set size in cylinders
- DDDEF name
- DSORG
- DSNTYPE (HFS, LIBRARY, PDS)
- Data set element type (LMOD, PNLENU, EXEC, etc.)
- Currently assigned logical volume
- Must be in master catalog (Y or N)
- Device type
- Currently assigned physical volume
- RECFM
- LRECL
- Whether data set is new in this order
- Whether rename is allowed (Yes, No, or Overidden)
- Data set type (Target, DLIB, or Operational)
- Volume sequence number (Tnn, Dnn, or Bnn)
- Subsystem Type (MVS, CICS, DB2, IMS, or NCP)
- Product, feature, or element name
- Special data set placement (FIRST or LAST)
- SMS-managed
- SMS-eligible
- SMS required
- Whether data set can be either a PDS or PDSE
- LPA required
- LPA eligible
- Link list eligible
- Volume Number (Tnn, Dnn, or Bnn)

There are new Change commands, too, and the Merge command now starts with the list you selected, which makes it much easier to merge like data sets. For instance, to merge all your US English panel data sets, select "Data Set Type," select "PNLENU" and "PNL," and issue the Merge line command...you're finished!

A note from the editors:

This article was published in our last issue, but the screen captures that illustrated its explanations escaped from the already fragile hold we have on reality, and did not as a result appear in the article. So, here it is again, this time complete and more usable.

Here is the first panel you will see when you select Option C from the **Modify System Layout options** panel. Here, you pick one attribute to work with. For example, you might pick "Element Type" to display certain kinds of data sets

```
CPPP605R ------------- Select Data Set View ( AM000321 ) -- Row 1 to 15 of 28
COMMAND ==>                                              SCROLL ==> HALF

Select a Data Set List View:

    Line Commands: (Select)

  S  Display          Data Set List Description
  -- --------------   -----------------------------------------------
     APF Required     APF Authorization Required (Yes or No)
     BLKSIZE          Data Set Block Size
     Current LVOL     Assigned Logical Volume
     Current PVOL     Assigned Physical Volume
     Cylinders        Current Data Set Size in Cylinders
     Data Set Type    DLIB, Target, Operational
     Device Type      Assigned Device Type
     DDDEF Name       Owning Product DDDEF Name
     DSNTYPE          HFS, LIBRARY or PDS
     DSORG            Data Set Organization
  s  Element type     Data Set Element Type (LMOD, PNLENU, EXEC, etc.)
  _  LNKLST Eligible  Eligible for placement in the Link List (Yes or No)
     LPA Eligible     Eligible for placement in LPA (Yes or No)
     LPA Required     Required in LPA List (Yes or No)
     LRECL            Logical Record Length
```

Next, you pick the values the dialog found for that attribute.

```
CPPP605S ------------ Select Values to Display ( AM000321 ) Row 19 to 30 of 37
COMMAND ==> list_                                       SCROLL ==> HALF

Select list for: Element type

Select values for which data sets are to be listed:

Primary Commands: (List)
    Line Commands: (Select SS)

    S   Values
    --  ------------------------------------------
        MSG
        MSGENP
        MSGENU
        MSGJPN
        PARM
    ss  PNL
        PNLENU
    ss  PNLJPN
        PROC
        PRODXML
        SAMP
        SAMPJPN
```

You can pick one, some, or all the values shown. Here, some of the data set types are displayed, and all the panels are being selected. Once you have made your choices, enter the List primary command to see the data set list.

Here's an example of this kind of list:

```
CPPP605U --------------- Data Set List ( AM000321 ) ------- Row 1 to 13 of 22
COMMAND ==> _                                           SCROLL ==> HALF

Data Set List for: Element type

Primary Commands: (? SET L F N P CHange OFile OList FC)
    Line Commands: (Merge eXpand Conflict Unmerge Attribs Space Resolve)

                                                          Physical
S Data Set Name                       Selected Value      Volume
- ------------------------------       --------------      --------
  SOMMVS.SGOSPNLS                      PNL                 JEKRS1
  SYS1.HRFPANL                         PNL                 JEKRS1
  SYS1.SBDTPNLO                        PNL                 JEKRS1
  SYS1.SBLSPNLO                        PNL                 JEKRS1
  SYS1.SISTPNLO                        PNL                 JEKRS2
  SYS1.SISTPNL1                        PNL                 JEKRS2
  SYS1.V2R1OMO.SHASPNLO                PNL                 JEKRS2
  SYS1.V2R1OMO.SIATPNLO                PNL                 JEKRS2
  TCPIP.SEZAPENU                       PNL                 JEKRS1
  AOP.SAOPPENU                         PNLENU              JEKRS1
  CPAC.ISPPLIB                         PNLENU              JEKRS1
  EOY.SEOYAENU                         PNLENU              JEKRS1
  GIM.SGIMPENU                         PNLENU              JEKRS1
```

This list shows all the panel libraries in this hypothetical order. This is how you can quickly get to lists from which the Merge command is very useful. If issued from this panel, it would show you a "merge candidate list" of all the eligible panel libraries.

Charting the changes — The Change and Merge Commands

People often ask us why, since IBM recommends no secondary space for link list libraries, ServerPac allocates them with secondary space by default. The real reason is, "Because that's how it's always been." But the reason we never changed was because some people build maintenance systems for which, unlike production systems, secondary extents are desirable.

Now, you can quickly set the secondary space to zero for all the link list data sets. The CH S(pace) command has a new operand—P0—that you can use to set secondary space to zero. Using *View and change*, select "Link list eligible," select "Yes," and then issue the new change command (CH S P0) to set them all to zero...and you're done! And, of course, secondary space is saved for each data set when you save a configuration for reuse with your next order.

There are other new change commands you can use to assign data sets directly to volumes (not "logical volumes," real volumes), change the high-level qualifiers of data sets, and—lots of people asked for this one—allow all data sets to be renamed.

Here are the other new change commands:
• You can use CH PVOL to assign target, DLIB, or operational data sets directly to a volume.
• CH DSN *HLQ* changes just the high-level qualifiers of data sets.
• CH RENAME command overrides the "not renameable" attribute for data sets.

You can still issue Change commands from any data set list panel, not only from the new *View and change*, panels. The Change command always starts with the list of data sets displayed, and shows you a list of candidate data sets from that list. Then you can exclude any you do not want to change and tell the dialog to apply the change to the rest of the data sets displayed.

The Merge command now starts with the data sets in the displayed list; it used to search the entire order for eligible data sets. When issued from the powerful *View and change*, option, we think you will find it much more useful. (However, if you liked the way it

used to work, you can just issue it from a list that includes all the data sets. For example, in *View and change*, select "link list eligible" and then select both "yes" and "no" and then issue the merge command against the desired merge target data set.)

Recommended System Layout — The Automatic Data Set Assignment story

Some time ago, a recommended system layout was published in *OS/390 Planning for Installation.* But it took a long time to install a ServerPac using that layout because the dialog did not lend itself well to that kind of structure. A spate of customer requirements (also called Requests) followed.

We are happy to report that with the ServerPac that becomes available at the same time as z/OS, a new function is introduced to help you get to the recommended system layout. Using this function is much faster than trying to create a layout manually. Most configurations take a minute or less to create.

The **Recommended System Layout** option will assign target and DLIB data sets to volumes for you according to a set of rules. But it won't move existing data sets unless you ask it to.

You can have the dialog assign all target and DLIB data sets, to create a layout for the first time; or, have it assign only the new ones, to preserve and existing layout. You can even have it reassign some of the old ones and all the new ones. (This "Partial" option is akin to compressing a PDS; sooner or later, you've got to "clean up.")

```
CPPP625B  -------- Automatic Data Set Assignment ( AM000321 ) -----------------
OPTION ==> A_


    A - ALL       Assign all target and DLIB data sets in the configuration
                  to physical volumes automatically.  This option creates a
                  recommended system layout.

    N - NEW       Add new data sets to an existing configuration.  This
                  option automatically assigns all new data sets, but
                  preserves the placement of previously-assigned data
                  sets in your saved configuration.

    P - PARTIAL   Assign new data sets and reassign some existing data sets
                  to physical volumes.  This option automatically assigns
                  all new data sets to physical volumes, as well as data
                  sets from selected volumes in the saved configuration.

  Default Device Type ==> 3390-3     (For example, 3390-3)
```

Once you pick All, New, or Partial, you see the current volume configuration.

If you did not bring forward a saved configuration, this is what you would see. There would be one target volume (MVSRES) and one DLIB volume (MVSDLB), both overallocated as you can see from the "Used & Reserved" column. On this panel, you can add volumes or change their attributes.

```
CPPP625C  -------- Automatic Data Set Assignment ( AM000321 )  Row 1 to 2 of 2
COMMAND ==> cr _                                          SCROLL ==> HALF

Current Volume Configuration                            Scope ==> ALL

Primary Commands: (? Reset CReate)
    Line Commands: (Select Insert List Move After Before eXclude)


    Phys.  Volume  Sequence  Device  Reserved  Used +    Volume     Existing
S   Volume Type    Number    Type    Space     Reserved  Threshold  Data
-   ------ ------  --------  ------  --------  --------  ---------  --------
    MVSRES TARGET  T01       3390-3  N         159 %     85 %       N
    MVSDLB DLIB    D01       3390-3  N         136 %     85 %       N
*********************************** Bottom of data ***********************************
```

When you're ready, you issue the Create command to start automatic data set assignment. When it has finished, the current volume configuration is redisplayed to show you what happened:

```
CPPP625C  -------- Automatic Data Set Assignment ( AM000321 )  Row 1 to 5 of 5
COMMAND ==> _                                             SCROLL ==> HALF

Current Volume Configuration                            Scope ==> ALL

Primary Commands: (? Reset CReate)
    Line Commands: (Select Insert List Move After Before eXclude)


    Phys.  Volume  Sequence  Device  Reserved  Used +    Volume     Existing
S   Volume Type    Number    Type    Space     Reserved  Threshold  Data
-   ------ ------  --------  ------  --------  --------  ---------  --------
    MVSRES TARGET  T01       3390-3  N         82 %      85 %       N
    TARG02 TARGET  T02       3390-3  N         63 %      85 %       N
    TARG03 TARGET  T03       3390-3  N         50 %      85 %       N
    MVSDLB DLIB    D01       3390-3  N         85 %      85 %       N
    DLIB02 DLIB    D02       3390-3  N         51 %      85 %       N
*********************************** Bottom of data ***********************************


        CPP0625014I Automatic Data Set Assignment Complete.
```

As you can see, this hypothetical order had its data sets assigned to four target and three DLIB volumes. The target data sets should be assigned to volumes in a way that is very close to the recommended system layout described in *z/OS Planning for Installation.*

With z/OS orders containing about 2,000 data sets, this is much faster than doing it by hand. This function leaves only the operational data sets for you to assign manually.

License Manager

There is a new component for z/OS, called IBM License Manager, or ILM. ILM will let you take advantage of workload pricing Charges. ILM requires

a more robust environment, but ServerPac will help you with most of it, particularly if you use the Full-System Replacement installation path.

ServerPac setup will focus on your first system by setting up an ILM Application Agent and License Server on that image.

ServerPac will set up these functions for ILM:
- RACF - All the profiles you need for ILM and the services it uses are created by a new installation job. With some review and editing to match your installation's standards, this job should help you get a running start. Also, it defines all the new high-level qualifiers in the order for you.
- Basic Sysplex - ServerPac's default is now a Basic Sysplex configuration, also called a Monoplex.
- System Logger - System Logger is set up for License Manager...and for OPERLOG and LOGREC, too.
- SMS - SMS is set up with the SMS classes needed to support System Logger, and with minimally-intrusive ACS routines that should not interfere with the allocation of other data sets.
- WLM - The system now comes up in Goal Mode using the default WLM configuration. (You can tailor the WLM configuration with its ISPF dialog.)
- TCP/IP - Loopback communication is enabled, so the Application Agent and Server can talk to each other. There is also a sample FTP configuration.
- OCSF - This z/OS element is entirely set up and ready to use.

All these things are either done when your order is built (ILM Application Agent and Server, Basic Sysplex, WLM, OCSF, TCP/IP, and SMS), or done by new installation jobs (RACF and Logger).

If you use the Software Upgrade installation path, you might want to generate these jobs for use as samples by selecting the Full-System Replacement installation path, use GENSKEL to create the jobs, copy them, and then go back to the Software Upgrade path to perform the installation. The jobs most useful as samples include:
- ALLOCDS, which allocates SMS, Couple, and ILM data sets
- RESTORE, which restores the SMS data sets
- RACFDRV and RACFTGT, which create needed RACF profiles
- DEFNLOGS, which creates the log streams.

When you have finished with the ServerPac installation jobs, all that will be left is finishing the FTP setup (after all, we don't know what hardware you have), installing the ILM workstation code, and installing and modifying the certificates.

Giving you what you asked for

We think you'll find ServerPac is better than ever.

If you have been involved in user groups like SHARE, or follow the IBM-MAIN discussion group on the Usenet, you will probably recognize several requirements in this article. This release of ServerPac includes a lot of things our customers have asked for over the past few years. Keep those requests and requirements coming! •

Our contributors

Debbie Beatrice — Debbie is a Senior Software Engineer in the user technologies area, focusing on installation ease of use. She has been with IBM for 19 years.

Ed Berkel — Ed joined IBM 29 years ago. He is a Senior Software Engineer and has worked in product development, performance test, and system test, In addition, he has been involved in Product Introduction programs for several releases of MVS and OS/390. He is currently on the z/OS Produuct Development Team, with responsibility fo technical marketing.

Franz-Peter Boley — Franz-Peter has been with IBM for eighteen years and has held a variety of software development positions within the IBM lab in Boeblingen, Germany.

Mary Ellen Coleman — Mary Ellen is an Advisory Software Engineer for the Information Development and Programming Technology team in Poughkeepsie, NY. She has been developing animation for WebSphere Application Server since 1997.

Paul DiMarzio — Paul is a Senior Software Engineer, with IBM since 1984 - all of it in z/OS and its predecessors (OS/390 and MVS). His focus for the past 10 years has been on the suitability of new application technologies to the z/OS environment. He has worked with UNIX, Windows NT, XML, and now enterprise Java and Web Services. He was the product manager for the introduction of WebSphere Application Server V4.0 to z/OS and now is on the WebSphere for z/OS technical marketing team.

Greg Dyck — Greg is a Senior Software Engineer in the z/OS development organization. He is responsible for the design and implementation of system functions in many areas of z/OS

John Eells — John is a former MVS and OS/390 system programmer, now a Senior Software Engineer in IBM Software Delivery and Fullfillment. SDF delivers all IBM's software and services packages including ServerPac, CBPDO, ESO, and SUF. His current assignments include ServerPac design.

Chris Hastings — Chris Hastings is a Senior Software Engineer at IBM's Poughkeepsie Lab. He has extensive experience in planning and writing documentation in support of large systems. His current assignments include being an advisor for the WebSphere for z/OS documentation, writing WebSphere for z/OS: Installation and Customization, and creating the help system for the WebSphere for z/OS Customization Dialog.

Barbara Klein — Barbara is the IMS Strategy, Plan and Brand Manager at IBM's Silicon Valley Laboratory, San Jose, California. Involved with IMS, CICS and DB2 for many years, she has held a variety of positions as Systems Engineer, Securities Industry Analyst, Telecommunications Product Planner, Systems Programmer, Business Planner, and manager of Programming Assurance and of IMS DC Design and Development.

Kirsten McDonald — Kirsten McDonald is a Software Engineer in the z/OS and OS/390 usability and Human Factors group. She has been with IBM for 13 years, and is currently working with the WebSphere for z/OS and e-business™ teams.

Robert Naumann — Robert is a student at the University of Nebraska in Omaha, studying towards a degree in Management of Information Systems at the Peter Keiwit Institute. Robert is currently working as a co-op in the User Technology Solutions department for IBM in Poughkeepsie, NY.

Kurt Quackenbush — Kurt Quackenbush is a Senior Software Engineer, and has been part of the SMP/E team and IBM for over 13 years. Currently he is a key designer and the development team leader for the SMP/E element of z/OS.

J.D. Ross — J.D. is a Staff Software Engineer at IBM Poughkeepsie and has worked on documentation for z/OS UNIX System Services and Linux for zSeries. He is also the webmaster of the Linux for zSeries website. Before joining IBM, J.D. served for several years as a system administrator, working with both the UNIX and Linux operating systems.

Stephen Saroka — Stephen works in the IBM Software Group and has been with IBM more than thirty years in architecture, design, development and servicing the IBM hardware and software mainframe systems. His current assignment in Electronic Strategy and Marketing Development organization are in e-delivery strategy for the S/390 and zSeries software service and products, and the marketing of IBM @server zSeries Service Update Facility and ShopzSeries worldwide.

James F. Steipp — James is an Advisory Software Engineer who has worked for IBM information development as a writer, planner, manager, Webmaster, and Wizard over the past 31 years. He currently spends most of his time on the Service Update Facility (SUF), the SoftCopy Librarian, and LookAt, with a little Web page development and z/OS wizardry thrown in for good measure.

Marilyn Zeppetelli — Marilyn is the z/OS Wizards development team leader. She developed the OS/390 TCPIP Configuration Assistant, OS/390 UNIX Configuration Assistant, and the OS/390 Planning Assistant for e-business. Her past assignments have included several years as a systems programmer, and several years in MVS Software Service (Level2) and VM Software Service (Level2).

Hot Topics Article Index

To help you more easily find articles that interest you, we provide the following topical index for the last four editions of the *Hot Topics Newsletter*. Below is a short summary of the contents of each of the four issues. Note that each issue contains a number of articles about a single topic as well as articles about individual topics. The last three editions also contain a supplement in which we deal with a single topic in greater depth.

OS/390 Hot Topics (August 1999), GA22-7431-00

Area of concentration: Language Environment and migration (8 articles)

Other topics: 10 articles

Total pages: 16

OS/390 Hot Topics (February 2000), GA22-7431-01

Area of concentration: TCP/IP (9 articles)

Other topics: 7 articles

Total pages: 32

Hot Topics Supplement (February 2000)

Topic: "UNIX on OS/390: A Powerful Blend"

Total pages: 16

OS/390 Hot Topics (August 2000), GA22-7431-02

Area of concentration: Dynamic Workload Balancing and Resource Management (16 articles)

Other topics: 12 articles

Total pages: 60

Hot Topics Supplement (August 2000)

Topic: "What the Heck is Parallel Sysplex Clustering Technology Anyway?"

Number of articles: 12

Total pages: 52

z/OS Hot Topics (February 2001), GA22-7501-00

Area of concentration: z/OS's new technologies (Intelligent Resource Director, etc.)

Number of articles: 21

Total pages: 50

z/OS-OS/390 Hot Topics Supplement (February 2001)

Topic: "Follow the Yellow Brick Information Superhighway"

(z/OS and OS/390 product information deliverables, wizards, LookAt, CD-ROMs, etc.)

Number of articles: 16

Total pages: 26

OS/390 Hot Topics (February 2001)

Note: this edition is a part of GA22-7501-00

Number of articles: 6

Total pages: 9

All editions are orderable in hardcopy, available on the CDROM collections, on the internet, and on the *zFavorites CD-ROM* contained in this issue. The URL for *Hot Topics* on the internet is
http://www-1.ibm.com/servers/s390/os390/bkserv/hot_topics.html

However, if you want to Bookmark the above URL and not risk typing the URL incorrectly, try the above mentioned *zFavorites CD-ROM*. It contains this URL and many other z-series URLs that you'll find useful.

Hot Topics Article Index By Topic

64-bit Architecture

"64-bit real operation in z/OS,"
OS/390 Hot Topics (February 2001), p. 6.

"64-bit real operation in OS/390 Release 10,"
z/OS Hot Topics (February 2001), p. 22.

ABTERMENC

"Attention LE Run-Time Options,"
OS/390 Hot Topics (February 2000), p. 20.

"Attention: LE Run-Time Options PR/SM,"
OS/390 Hot Topics (August 2000), p. 50.

Animations of product concepts

"We're animated!,"
z/OS-OS/390 Hot Topics Supplement (February 2001), p. 21.

BookManager

"CD Collections to contain 'Try and Buy' for Build/NT,"
OS/390 Hot Topics (August 1999), p. 14.

"The bookworm's workhorse – IBM BookManager: the complete online documentation solution,"
z/OS-OS/390 Hot Topics Supplement (February 2001), p. 15.

"Easy on your eye – IBM SoftCopy Reader: the new Java-based viewer for softcopy,"
z/OS-OS/390 Hot Topics Supplement (February 2001), p. 17.

Books, ordering (see *Ordering*)

Build/NT

"CD Collections to contain 'Try and Buy' for Build/NT," *OS/390 Hot Topics (August 1999)*, p. 14.

C/C++

"Using XPLINK to improve C/C++ program performance,"
OS/390 Hot Topics (August 2000), p. 50.

"So, what gets XPLINK started? (LE, C/C++),"
OS/390 Hot Topics (August 2000), p. 51.

Stepping to z/OS… continued from page 44

Manager (ILM). This involves preparations across three distinct areas:
a. ILM infrastructure requirements
b. LPAR capacity planning
c. License certificate management

You will undoubtedly want to prepare for the License Manager as much as possible before your z/OS system first meets the z900. The infrastructure preparation involves setting up TCP/IP, Unix System Services, the Open Cryptographic Support Facility, some security definitions, logstreams managed by System Logger, etc. This setup can be accomplished on any OS/390 R8, R9, or R10 system as well as on z/OS itself on any supported hardware.

LPAR Capacity planning can be done from any system.

Management of license certificates is the one function that must run from a z/OS system on a zSeries server.

In summary, many accounts will find that OS/390 V2R10 is a valuable step on the way to running z/OS. Although it is not a required step, R10 provides the opportunity or reassurance of testing Assembler language products when they first encounter 64-bit reall addresses. It also provides an opportunity for configuring the License Manager prerequisites before your z900 first encounters the z/OS operating system. This article provides just a short introduction to these considerations for hardware and software migration. Please refer to Chapter 6 of the *z/OS Planning for Installation* manual for a more thorough treatment of the topic. •

Stepping to z/OS

ED BERKEL

Some customers wonder if they should install OS/390 R10 prior to installing z/OS. There isn't a simple answer valid for all customers. While going to R10 is NOT required, it may be desirable under certain circumstances. Therefore, as with many other situations, the answer for you depends on your own unique requirements and configuration.

Diane Attebury

First, let's cover some of the easy answers.
- If you already have plans to install OS/390 R10 this year, continue with them.
- If you currently run a release of OS/390 that is no longer in service, and therefore outside the coexistence policy allowing resources to be shared with z/OS, you would be wise to contract with IBM Global Services to pursue the feasibility of a migration to OS/390 R10.
- Similarly, if you are currently running OS/390 R6 or R7 and considering a migration to z/OS V1R2 (or higher), you would be outside the coexistence policy. Therefore, you would be wise to migrate to either OS/390 R10 or z/OS 1.1. See the later part of this article to help choose which is better for you. From either of them, you can migrate to z/OS V1R2 (or higher) when you are ready.
- If you currently run OS/390 on a processor that is not supported by z/OS, but is supported by

R10, you might find it worthwhile to upgrade to R10. Remember that OS/390 R10 will remain orderable until March 2002. By upgrading to OS/390 R10, you extend your options for future upgrades to z/OS. That is, the normal z/OS coexistence policy would support your OS/390 R10 system to share resources with z/OS R4, so taking the step to OS/390 R10 could help you have a smoother migration later. If you do decide to upgrade your processor and want to get to z/OS quickly, recall that the *z/OS V1R1 Upgrade Package for OS/390 V2R10* will be orderable until at least March, 2002 as well. So you can quickly upgrade to z/OS V1R1 if you need to.

You will want to take into account that z/OS has two key requirements when z/OS runs on a zSeries server. You will want to consider these requirements as part of your hardware and software migration plans over the next two years. These considerations apply whether the first place your programs or vendor products meet a z900 is on your production system, or your test system, or your disaster recovery site. So the hardware available at (or planned for) your production, test, and disaster recovery locations should be assessed when considering your software migrations. Specifically:

1. When running on a z900, z/OS requires z/Architecture mode with its 64-bit real addressing.

If you have written your own Assembler language programs that depend on real storage addresses, or have vendor products with those dependencies, you might prefer to have R10 as the operating system

you use when the first z900 is installed, rather than having z/OS active then. That's because with R10 you can switch between ESA/390 (31-bit) and z/Architecture (64-bit) real addressing as you test out your own programs or your vendor programs. You cannot switch addressing modes with z/OS. Many accounts might find that this testing consideration alone is sufficient reason to install OS/390 R10 now, and begin running a 64-bit LPAR as soon as you have access to a z900.

Remember there shouldn't be any impact to normal applications or products. We are discussing 64-bit real addresses, not the virtual addresses used by most applications.

The earlier paragraph stated that this testing consideration applies for the first z900 installed. Assuming you have similar software stacks or workloads, your testing can be completed with OS/390 R10 on that first z900. Then you can upgrade to z/OS and propagate z/OS prior to upgrading the second or subsequent z900 in your enterprise.

2. When running on a z900, z/OS is priced according to Workload License Charges.

After the IBM License Manager pre-production period, you will need to activate the IBM License

GA22-7501-01